Journal of Machine Learning Research XX (2017) 1-25 Submitted 01/2017; Published XX/XX

Patchwork Kriging for Large-scale Gaussian Process Regression

Chiwoo Park CPARK5@FSU.EDU
Department of Industrial and Manufacturing Engineering
Florida State University
2525 Pottsdamer St., Tallahassee, FL 32310-6046, USA.

Daniel Apley APLEY@NORTHWESTERN.EDU
Dept. of Industrial Engineering and Management Sciences
Northwestern University
2145 Sheridan Rd., Evanston, IL 60208-3119, USA.

Editor:

arXiv:1701.06655v1 [cs.LG] 23 Jan 2017

Abstract

This paper presents a new approach for Gaussian process (GP) regression for large datasets. The approach involves partitioning the regression input domain into multiple local regions with a different local GP model fitted in each region. Unlike existing local partitioned GP approaches, we introduce a technique for patching together the local GP models nearly seamlessly to ensure that the local GP models for two neighboring regions produce nearly the same response prediction and prediction error variance on the boundary between the two regions. This effectively solves the well-known discontinuity problem that degrades the boundary accuracy of existing local partitioned GP methods. Our main innovation is to represent the continuity conditions as additional pseudo-observations that the differences between neighboring GP responses are identically zero at an appropriately chosen set of boundary input locations. To predict the response at any input location, we simply augment the actual response observations with the pseudo-observations and apply standard GP prediction methods to the augmented data. In contrast to heuristic continuity adjustments, this has an advantage of working within a formal GP framework, so that the GP-based predictive uncertainty quantification remains valid. Our approach also inherits a sparse block-like structure for the sample covariance matrix, which results in computationally efficient closed-form expressions for the predictive mean and variance. In addition, we provide a new spatial partitioning scheme based on a recursive space partitioning along local principal component directions, which makes the proposed approach applicable for regression domains having more than two dimensions. Using three spatial datasets and two high dimensional datasets, we investigate the numerical performance of the approach and compare it to several state-of-the-art approaches.

Keywords: Local Kriging, Model Split and Merge, Pseudo Observations, Spatial Partition

1. Introduction

Gaussian process (GP) regression is a popular Bayesian nonparametric approach for nonlinear regression (Rasmussen and Williams, 2005). A GP prior is assumed for the unknown regression function, and the posterior estimate of the function is from this prior, combined

©2017 Chiwoo Park and Daniel Apley.

with noisy (or noiseless, for deterministic simulation response surfaces) response observations. The posterior estimate can be easily derived in a simple closed form using the properties induced by the GP prior, and the estimator has several desirable properties, e.g., it is the best linear unbiased estimator under the assumed model and offers convenient quantification of the prediction error uncertainty. Its conceptual simplicity and attractive properties are major reasons for its popularity. On the other hand, the computational expense for evaluating the closed form solution is proportional to N^3, where N denotes the number of observations, which can be prohibitively expensive for large N. Broadly speaking, this paper concerns fast computation of the GP regression estimate for a large N.

The major computational bottleneck for GP regression is the inversion of a $N \times N$ sample covariance matrix, which is also often poorly numerically conditioned. Different approaches for representing or approximating the sample covariance matrix with a more efficiently invertible form have been proposed. The approaches can be roughly categorized as sparse approximations, low-rank approximations, or local approximations. Sparse methods represent the sample covariance with a sparse version, e.g. by applying a covariance tapering technique (Furrer et al., 2006; Kaufman et al., 2008), using a compactly supported covariance function (Gneiting, 2002), or using a Gaussian Markov approximation of a GP (Lindgren et al., 2011). The inversion of a sparse matrix is less computationally expensive than the inversion of a non-sparse matrix of the same size. Low-rank approximations introduces latent variables and assume a certain independence conditioned on the latent variables (Seeger et al., 2003; Snelson and Ghahramani, 2006), so that the resulting sample covariance matrix has reduced rank. The (pseudo)inversion of a $N \times N$ matrix of rank M can be computed with reduced $O(NM^2)$ expense. One class of local approximations partition the input domain into a set of local regions and assume an independent GP regression model within each region. The resulting sample covariance matrix is a block diagonal matrix of local sample covariance matrices, and inverting the block diagonal matrix is much cheaper computationally.

Such local approximation approaches have many advantages. By their local nature, they adapt better to local/nonstationary data features, and independent local approximation models can be computed in parallel to reduce total computation time. Their major weakness is that two local models for two neighboring local regions produce different predictions at the boundary between the regions, resulting in boundary discontinuity for regression predictive function. This boundary discontinuity is not just an aesthetic problem, as it was empirically shown that greater discontinuity implies greater degradation in prediction accuracy, particularly around the boundaries of the local regions (Park and Huang, 2016). This discontinuity issue has been addressed in different ways. Perhaps the most popular approach is to smooth out some of the discontinuity by using some weighted average across the local models or across multiple sets of local models via a Dirichlet mixture (Rasmussen and Ghahramani, 2002), a treed mixture (Gramacy and Lee, 2008), Bayesian model averaging (Tresp, 2000; Chen and Ren, 2009; Deisenroth and Ng, 2015), or locally weighted projections (Nguyen-Tuong et al., 2009). Other, related approaches use an additive covariance functions consisting of a global covariance and a local covariance (Snelson and Ghahramani, 2007; Vanhatalo and Vehtari, 2008), construct a local model for each testing point (Gramacy and Apley, 2014), or use a local partition but constrain the local models for continuity (Park et al., 2011; Park and Huang, 2016).

In this work we use a partitioned input domain like Park and Huang (2016), but we introduce a different form of continuity constraints that are more easily and more naturally integrated into the GP modeling framework. In Park and Huang (2016), the local GP models for neighboring regions are constrained to have the same predictive means on the boundaries of the local regions, but not necessarily the same predictive variances. Also, the values of the predictive means on the boundary must be explicitly specified, which involves solving a large linear system of equations. In contrast, we propose a new approach to impose continuity constraints for the local GP models on the boundaries. We consider a set of GPs that are defined as the differences between the responses for the local GPs in neighboring regions. Continuity implies that the these differenced GPs are identically zero along the boundary between neighboring regions. Hence, we impose continuity constraints by treating the values of the differenced GPs at a specified set of boundary points as all having been observed" to be zero, and we refer to these zero-valued differences as pseudo-observations. We can then conveniently incorporate continuity constraints by simply augmenting the actual set of response observations with the set of pseudo-observations, and then using standard GP modeling to calculate the posterior predictive distribution given the augmented set of observations. This pseudo-observation approach enforces that the local models for neighboring regions produce the same values for both the predictive means and and the prediction error variances at the boundary points between the local regions. We note that *observing* the differenced GPs to be zero at a set of boundary points is essentially equivalent to *assuming* continuity at these points without imposing any further assumptions on the nature of the GPs. Unlike the approach of Park and Huang (2016), observing the pseudo-observations of the differenced GPs to be zero does not require specifying the actual values of the predictive means and variances on the boundary. We also demonstrate that enforcing the local models to have the same predictive means and variances on the boundary improves the accuracy of the local predictive variance estimates, as well as the local predictive mean estimates. Moreover, our new approach is applicable for high dimensional regression problems, while Park and Huang (2016) is only applicable for one- or two-dimensional problems. We view our approach as patching together a collection of local GP regression models using the boundary points as stitches, and, hence, we refer to it as *patchwork kriging*.

The remainder of the paper is organized as follows. Section 2 briefly reviews the general GP regression problem and notational convention. Section 3 presents the core methodology of the patchwork kriging approach, including the prior model assumptions, the pseudo-observation definition, the resulting posterior predictive mean and variance equations, and the various properties of the estimates with relevant proofs and illustrations. Section 4 includes discussions on several relevant model selection issues, including spatial partitioning and covariance parameter estimation. Section 5 investigates the numerical performance of the proposed method and compares it with the exact GP regression (i.e., the GP regression without partitions, using the entire data set to predict each point) and with (Park and Huang, 2016) and other state-of-the-art methods using five real data sets. Finally, Section 6 concludes the paper with discussion.

2. Gaussian Process Regression

Consider the general regression problem of estimating an unknown predictive function f that relates a d dimensional predictor $x \in \mathbb{R}^d$ to a real response y, using noisy observations $\mathcal{D} = \{(x_i, y_i), i = 1, \ldots, N\}$,

$$y_i = \mu + f(x_i) + \epsilon_i, \qquad i = 1, \ldots, N,$$

where $\epsilon_i \sim \mathcal{N}(0, \sigma^2)$ is white noise, independent of $f(x_i)$. We assume that $\mu = 0$. Otherwise, one can normalize y_i by subtracting the sample mean of y_i's from y_i. Notice that we do not use bold font for the multivariate predictor x_i and reserve bold font for the collection of observed predictor locations, $\mathbf{x} = [x_1, x_2, \ldots, x_N]^T$.

In a GP regression for this problem, one assumes that f is a realization of a zero-mean Gaussian process having covariance function $c(\cdot, \cdot)$ and then uses the observations $\mathcal{D}$ to obtain the posterior predictive distribution of f at an arbitrary x_*, denoted by $f_* = f(x_*)$. Denote $\boldsymbol{y} = [y_1, y_2, \ldots, y_N]^T$. The joint distribution of $(f_*, \boldsymbol{y})$ is

$$P(f_*, \boldsymbol{y}) = \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} c_{**} & \boldsymbol{c}_{\mathbf{x}*}^T \\ \boldsymbol{c}_{\mathbf{x}*} & \sigma^2 \boldsymbol{I} + \boldsymbol{C}_{\mathbf{xx}} \end{bmatrix}\right),$$

where $c_{**} = c(x_*, x_*)$, $\boldsymbol{c}_{\mathbf{x}*} = (c(x_1, x_*), \ldots, c(x_N, x_*))^T$ and $\boldsymbol{C}_{\mathbf{xx}}$ is an $N \times N$ matrix with $(i, j)^{th}$ entry $c(x_i, x_j)$. The subscripts on $c_{**}, \boldsymbol{c}_{\mathbf{x}*}$, and $\boldsymbol{C}_{\mathbf{xx}}$ indicate the two sets of locations between which the covariance is computed, and we have abbreviated the subscript x_* as $*$. Applying the Gaussian conditioning formula to the joint distribution gives the predictive distribution of f_* given $\boldsymbol{y}$ (Rasmussen and Williams, 2005),

$$P(f_*|\boldsymbol{y}) = \mathcal{N}(\boldsymbol{c}_{\mathbf{x}*}^T(\sigma^2 \boldsymbol{I} + \boldsymbol{C}_{\mathbf{xx}})^{-1}\boldsymbol{y}, c_{**} - \boldsymbol{c}_{\mathbf{x}*}^T(\sigma^2 \boldsymbol{I} + \boldsymbol{C}_{\mathbf{xx}})^{-1}\boldsymbol{c}_{\mathbf{x}*}). \tag{1}$$

The predictive mean $\boldsymbol{c}_{\mathbf{x}*}^T(\sigma^2 \boldsymbol{I} + \boldsymbol{C}_{\mathbf{xx}})^{-1}\boldsymbol{y}$ is taken to be the point prediction of $f(x)$ at location x_*, and its uncertainty is measured by the predictive variance $c_{**} - \boldsymbol{c}_{\mathbf{x}*}^T(\sigma^2 \boldsymbol{I} + \boldsymbol{C}_{\mathbf{xx}})^{-1}\boldsymbol{c}_{\mathbf{x}*}$. Efficient calculation of the predictive mean and variance for large data sets has been the focus of much research.

3. Patchwork Kriging

As mentioned in the introduction, for efficient computation we replace the GP regression by a set of local GP models on some partition of the input domain, in a manner that enforces a level of continuity in the local GP model responses over the boundaries separating their respective regions. To specify the idea more precisely, consider a spatial partition of the input domain of $f(x)$ into M local regions $\{\Omega_k : k = 1, 2, ..., K\}$, and define $f_k(x)$ as the local GP approximation of $f(x)$ at $x \in \Omega_k$. Temporarily ignoring the continuity requirements, the local models are assumed to follow independent GP priors:

Assumption 1 *Each $f_k(x)$ follows a GP prior distribution with zero mean and covariance function $c_k(\cdot, \cdot)$, and the $f_k(x)$'s are mutually independent a priori.*

Partition the training set $\mathcal{D}$ into $\mathcal{D}_k = \{(x_i, y_i) : x_i \in \Omega_k\}$ $(k = 1, 2, ..., K)$, and denote $\mathbf{x}_k = \{x_i : x_i \in \Omega_k\}$ and $\boldsymbol{y}_k = \{y_i : x_i \in \Omega_k\}$. By the independence part of Assumption 1,

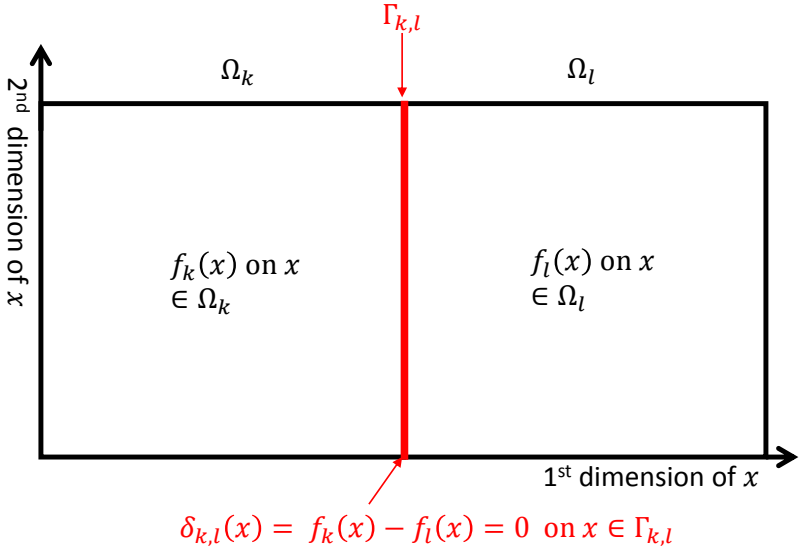
$\Gamma_{k,l}$
Ω_k
Ω_l
2nd dimension of x
$f_k(x)$ on x $\in \Omega_k$
$f_l(x)$ on x $\in \Omega_l$
1st dimension of x
$\delta_{k,l}(x) = f_k(x) - f_l(x) = 0$ on $x \in \Gamma_{k,l}$

Figure 1: Illustration of the notation and concepts for defining the local models $f_k(x)$ and $f_l(x)$. Ω_k and Ω_l represent two local regions resulting from some appropriate spatial partition (discussed later) of the regression input domain. The (posterior distributions for the) GP functions $f_k(x)$ and $f_l(x)$ represent the local approximations of the regression function on Ω_k and Ω_l, respectively. The subset $\Gamma_{k,l}$ represents the interfacial boundary between the two regions Ω_k and Ω_l, and $\delta_{k,l}(x)$ is defined as the difference between $f_k(x)$ and $f_l(x)$, which is identically zero on $\Gamma_{k,l}$ by the continuity assumption.

the predictive distribution of $f_k(x)$ given $\mathcal{D}$ is equivalent to the predictive distribution of $f_k(x)$ given $\mathcal{D}_k$, which gives the standard local GP solution with no continuity requirements. The primary problem with this formulation is that the predictive distributions of $f_k(x)$ and $f_l(x)$ are not equal on the boundary of their neighboring regions Ω_k and Ω_l.

Our objective is to improve the local kriging prediction by enforcing $f_k(x) = f_l(x)$ on their shared boundary. The key idea is illustrated in Figure 1 and described as follows. For two neighboring regions Ω_k and Ω_l, let $\Gamma_{k,l} = \overline{\Omega}_k \cap \overline{\Omega}_l$ denote their shared boundary, where $\overline{\Omega}_k$ is the closure of Ω_k. Instead of attempting to modify the *a priori* independence part of Assumption 1 to enforce boundary continuity (which would not be straightforward and could lead to computational intractability), we propose a much simpler continuity correction. For each pair of neighboring regions Ω_k and Ω_l, we define the auxiliary process $\delta_{k,l}(x)$ to be the difference between the two local GP models,

$$\delta_{k,l}(x) = f_k(x) - f_l(x) \text{ for } x \in \Gamma_{k,l}, \tag{2}$$

and it is only defined for $k < l$ to avoid any duplicated definition of the auxiliary process. By the definition and under Assumption 1, $\delta_{k,l}(x)$ is a Gaussian process with zero mean

and covariance function $c_k(\cdot,\cdot) + c_l(\cdot,\cdot)$, and its covariance with $f_j(x)$ is

$$\mathbb{C}\mathrm{ov}(\delta_{k,l}(x_1), f_j(x_2)) = \begin{cases} c_k(x_1,x_2) & \text{if } k=j \\ -c_l(x_1,x_2) & \text{if } l=j \\ 0 & \text{otherwise.} \end{cases} \tag{3}$$

Likewise, $\delta_{k,l}(x)$ and $\delta_{u,v}(x)$ are correlated with covariance

$$\mathbb{C}\mathrm{ov}(\delta_{k,l}(x_1), \delta_{u,v}(x_2)) = \begin{cases} c_k(x_1,x_2) & \text{if } k=u, l\neq v \\ c_l(x_1,x_2) & \text{if } l=v, k\neq u \\ -c_k(x_1,x_2) & \text{if } k=v, l\neq u \\ -c_l(x_1,x_2) & \text{if } l=u, k\neq v \\ c_k(x_1,x_2)+c_l(x_1,x_2) & \text{if } k=u, l=v \\ 0 & \text{otherwise.} \end{cases} \tag{4}$$

Our strategy for enforcing boundary continuity between $f_k(x)$ and $f_l(x)$ is to augment the observed data $\mathcal{D}$ with the pseudo observations $\delta_{k,l}(x) = 0$ for a finite number of input locations $x \in \Gamma_{k,l}$. The choice of the boundary input locations will be discussed later. Notice that observing" $\delta_{k,l}(x) = 0$ is equivalent to observing that $f_k(x) = f_l(x)$ without observing the actual values of $f_k(x)$ and $f_l(x)$. Thus, if we augment $\mathcal{D}$ to include these pseudo observations when calculating the posterior predictive distributions of $f_k(x)$ and $f_l(x)$, it will force the posterior distributions of $f_k(x)$ and $f_l(x)$ to be the same at each boundary input location x, because observing $\delta_{k,l}(x) = 0$ means that we have observed $f_k(x)$ and $f_l(x)$ to be the same (see Theorem 1, below, for a formal proof). This implies that their posterior means (which are the GP regression predictive functions) and their posterior variances (which quantifies the uncertainty in the predictions) will both be equal.

Let $\boldsymbol{x}_{k,l}$ denote the set of B input boundary locations chosen in $\Gamma_{k,l}$, let $\boldsymbol{\delta}_{k,l}$ denote the set of associated $\delta_{k,l}(x)$ values, and let $\boldsymbol{\delta}$ denote the collection of all $\boldsymbol{\delta}_{k,l}$'s. Note that the observed pseudo value of $\boldsymbol{\delta}$ will be a vector of zeros, but its prior distribution (prior to observing the pseudo values) is represented by the above covariance expressions. Additionally, let $f_*^{(k)} = f_k(x_*)$ denote the value of the local approximation $f_k(x)$ at any $x_* \in \Omega_k$, and let $\boldsymbol{y}^T = (\boldsymbol{y}_1^T, \boldsymbol{y}_2^T, \ldots, \boldsymbol{y}_K^T)$. The prior joint distribution of $f_*^{(k)}$, $\boldsymbol{y}$ and $\boldsymbol{\delta}$ is

$$\begin{bmatrix} f_*^{(k)} \\ \boldsymbol{y} \\ \boldsymbol{\delta} \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} 0 \\ \boldsymbol{0} \\ \boldsymbol{0} \end{bmatrix}, \begin{bmatrix} c_{**} & \boldsymbol{c}_{*\mathcal{D}}^{(k)} & \boldsymbol{c}_{*,\delta}^{(k)} \\ \boldsymbol{c}_{\mathcal{D}*}^{(k)} & \boldsymbol{C}_{\mathcal{D}\mathcal{D}} & \boldsymbol{C}_{\mathcal{D},\delta} \\ \boldsymbol{c}_{\delta,*}^{(k)} & \boldsymbol{C}_{\delta,\mathcal{D}} & \boldsymbol{C}_{\delta,\delta} \end{bmatrix} \right),$$

where $\boldsymbol{c}_{\mathcal{D}*}^{(k)}$ is the blocked column vector of K column vectors with the uth column vector $\mathbb{C}\mathrm{ov}(\boldsymbol{y}_u, f_*^{(k)})$, $\boldsymbol{c}_{\delta*}^{(k)}$ is the blocked column vector of $K(K-1)/2$ column vectors with the $((u-1)K+v)$th column vector $\mathbb{C}\mathrm{ov}(\boldsymbol{\delta}_{u,v}, f_*^{(k)})$, $\boldsymbol{C}_{\mathcal{D}\mathcal{D}}$ is the block diagonal matrix of $K(K-1)/2$ diagonal blocks with the $((u-1)K+v)$th diagonal block $\mathbb{C}\mathrm{ov}(\boldsymbol{y}_u, \boldsymbol{y}_v)$, $\boldsymbol{C}_{\delta,\mathcal{D}}$ is the block matrix of $K(K-1)/2 \times K$ blocks with its $((u-1)K+v, l)$th block $\mathbb{C}\mathrm{ov}(\boldsymbol{\delta}_{u,v}, \boldsymbol{y}_l)$, and $\boldsymbol{C}_{\delta,\delta}$ is the block matrix of $K(K-1)/2 \times K(K-1)/2$ blocks with its $((u-1)K+v, (l-1)*K+m)$th

block $\mathbb{C}\mathrm{ov}(\boldsymbol{\delta}_{u,v}, \boldsymbol{\delta}_{l,m})$. Note that joint covariance matrix is very sparse, because of the many zero values in (3) and (4).

From the standard GP modeling results applied to the augmented data, the posterior predictive distribution of $f_*^{(k)}$ given $\boldsymbol{y}$ and the pseudo observations $\boldsymbol{\delta}$ is Gaussian distribution with mean

$$\mathbb{E}[f_*^{(k)}|\boldsymbol{y}, \boldsymbol{\delta}] = [\boldsymbol{c}_{*\mathcal{D}}^{(k)}, \boldsymbol{c}_{*,\delta}^{(k)}] \begin{bmatrix} \boldsymbol{C}_{\mathcal{D}\mathcal{D}} & \boldsymbol{C}_{\mathcal{D},\delta} \\ \boldsymbol{C}_{\delta,\mathcal{D}} & \boldsymbol{C}_{\delta,\delta} \end{bmatrix}^{-1} \begin{bmatrix} \boldsymbol{y} \\ \boldsymbol{\delta} \end{bmatrix}.$$

Using the partitioned matrix inversion formula

$$\begin{bmatrix} \boldsymbol{A} & \boldsymbol{B} \\ \boldsymbol{B}^T & \boldsymbol{D} \end{bmatrix}^{-1} = \begin{bmatrix} (\boldsymbol{A} - \boldsymbol{B}\boldsymbol{D}^{-1}\boldsymbol{B}^T)^{-1} & -(\boldsymbol{A} - \boldsymbol{B}\boldsymbol{D}^{-1}\boldsymbol{B}^T)^{-1}\boldsymbol{B}\boldsymbol{D}^{-1} \\ -\boldsymbol{D}^{-1}\boldsymbol{B}^T(\boldsymbol{A} - \boldsymbol{B}\boldsymbol{D}^{-1}\boldsymbol{B}^T)^{-1} & (\boldsymbol{D} - \boldsymbol{B}^T\boldsymbol{A}^{-1}\boldsymbol{B})^{-1} \end{bmatrix}, \quad (5)$$

we have

$$\begin{aligned}\mathbb{E}[f_*^{(k)}|\boldsymbol{y}, \boldsymbol{\delta}] =& \boldsymbol{c}_{*\mathcal{D}}^{(k)}(\boldsymbol{C}_{\mathcal{D}\mathcal{D}} - \boldsymbol{C}_{\mathcal{D}\delta}\boldsymbol{C}_{\delta,\delta}^{-1}\boldsymbol{C}_{\delta\mathcal{D}}^T)^{-1}\boldsymbol{y} - \boldsymbol{c}_{*,\delta}^{(k)}\boldsymbol{C}_{\delta,\delta}^{-1}\boldsymbol{C}_{\mathcal{D}\mathcal{D}}^T(\boldsymbol{C}_{\mathcal{D}\mathcal{D}} - \boldsymbol{C}_{\mathcal{D}\delta}\boldsymbol{C}_{\delta,\delta}^{-1}\boldsymbol{C}_{\mathcal{D}\delta}^T)^{-1}\boldsymbol{y} \\ & - \boldsymbol{c}_{*\mathcal{D}}^{(k)}(\boldsymbol{C}_{\mathcal{D}\mathcal{D}} - \boldsymbol{C}_{\mathcal{D}\delta}\boldsymbol{C}_{\delta,\delta}^{-1}\boldsymbol{C}_{\mathcal{D}\delta}^T)^{-1}\boldsymbol{C}_{\mathcal{D}\mathcal{D}}\boldsymbol{C}_{\delta,\delta}^{-1}\boldsymbol{\delta} + \boldsymbol{c}_{*,\delta}^{(k)}(\boldsymbol{C}_{\delta,\delta} - \boldsymbol{C}_{\mathcal{D}\delta}^T\boldsymbol{C}_{\mathcal{D}\mathcal{D}}^{-1}\boldsymbol{C}_{\mathcal{D}\delta})^{-1}\boldsymbol{\delta}.\end{aligned}$$

In particular, for $\boldsymbol{\delta} = \mathbf{0}$, which is the only value for the pseudo observations that we need, we have

$$\mathbb{E}[f_*^{(k)}|\boldsymbol{y}, \boldsymbol{\delta} = \mathbf{0}] = (\boldsymbol{c}_{*\mathcal{D}}^{(k)} - \boldsymbol{c}_{*\delta}^{(k)}\boldsymbol{C}_{\delta,\delta}^{-1}\boldsymbol{C}_{\mathcal{D}\delta}^T)(\boldsymbol{C}_{\mathcal{D}\mathcal{D}} - \boldsymbol{C}_{\mathcal{D}\delta}\boldsymbol{C}_{\delta,\delta}^{-1}\boldsymbol{C}_{\mathcal{D}\delta}^T)^{-1}\boldsymbol{y}. \quad (6)$$

Similarly, the posterior predictive variance of $f_*^{(k)}$ given $\boldsymbol{y}$ and $\boldsymbol{\delta}$ is

$$\mathbb{V}\mathrm{ar}[f_*^{(k)}|\boldsymbol{y}, \boldsymbol{\delta}] = c_{**} - [\boldsymbol{c}_{*\mathcal{D}}^{(k)}, \boldsymbol{c}_{*,\delta}^{(k)}] \begin{bmatrix} \boldsymbol{C}_{\mathcal{D}\mathcal{D}} & \boldsymbol{C}_{\mathcal{D},\delta} \\ \boldsymbol{C}_{\delta,\mathcal{D}} & \boldsymbol{C}_{\delta,\delta} \end{bmatrix}^{-1} \begin{bmatrix} \boldsymbol{c}_{\mathcal{D}*}^{(k)} \\ \boldsymbol{c}_{\delta*}^{(k)} \end{bmatrix}.$$

Applying the matrix inversion result (5) to this variance expression, we have

$$\begin{aligned}\mathbb{V}\mathrm{ar}[f_*^{(k)}|\boldsymbol{y}, \boldsymbol{\delta}] =& c_{**} - \boldsymbol{c}_{*\mathcal{D}}^{(k)}(\boldsymbol{C}_{\mathcal{D}\mathcal{D}} - \boldsymbol{C}_{\mathcal{D}\delta}\boldsymbol{C}_{\delta,\delta}^{-1}\boldsymbol{C}_{\delta\mathcal{D}}^T)^{-1}\boldsymbol{c}_{\mathcal{D}*}^{(k)} \\ & - \boldsymbol{c}_{*,\delta}^{(k)}\boldsymbol{C}_{\delta,\delta}^{-1}\boldsymbol{C}_{\mathcal{D}\mathcal{D}}^T(\boldsymbol{C}_{\mathcal{D}\mathcal{D}} - \boldsymbol{C}_{\mathcal{D}\delta}\boldsymbol{C}_{\delta,\delta}^{-1}\boldsymbol{C}_{\mathcal{D}\delta}^T)^{-1}\boldsymbol{c}_{\mathcal{D}*}^{(k)} \\ & - \boldsymbol{c}_{*\mathcal{D}}^{(k)}(\boldsymbol{C}_{\mathcal{D}\mathcal{D}} - \boldsymbol{C}_{\mathcal{D}\delta}\boldsymbol{C}_{\delta,\delta}^{-1}\boldsymbol{C}_{\mathcal{D}\delta}^T)^{-1}\boldsymbol{C}_{\mathcal{D}\mathcal{D}}\boldsymbol{C}_{\delta,\delta}^{-1}\boldsymbol{c}_{\delta*}^{(k)} \\ & + \boldsymbol{c}_{*,\delta}^{(k)}(\boldsymbol{C}_{\delta,\delta} - \boldsymbol{C}_{\mathcal{D}\delta}^T\boldsymbol{C}_{\mathcal{D}\mathcal{D}}^{-1}\boldsymbol{C}_{\mathcal{D}\delta})^{-1}\boldsymbol{c}_{\delta*}^{(k)} \\ =& c_{**} - \boldsymbol{c}_{*,\delta}^{(k)}\boldsymbol{C}_{\delta,\delta}^{-1}\boldsymbol{c}_{\delta*}^{(k)} \\ & - (\boldsymbol{c}_{*\mathcal{D}}^{(k)} - \boldsymbol{c}_{*\delta}^{(k)}\boldsymbol{C}_{\delta,\delta}^{-1}\boldsymbol{C}_{\mathcal{D}\delta}^T)(\boldsymbol{C}_{\mathcal{D}\mathcal{D}} - \boldsymbol{C}_{\mathcal{D}\delta}\boldsymbol{C}_{\delta,\delta}^{-1}\boldsymbol{C}_{\mathcal{D}\delta}^T)^{-1}(\boldsymbol{c}_{\mathcal{D}*}^{(k)} - \boldsymbol{C}_{\mathcal{D}\delta}\boldsymbol{C}_{\delta,\delta}^{-1}\boldsymbol{c}_{\delta*}^{(k)}).\end{aligned}$$

The specific form of continuity imposed by including the pseudo observations $\boldsymbol{\delta} = \mathbf{0}$ is that the posterior predictive means and variances of $f_*^{(k)}$ and $f_*^{(l)}$ for two neighboring regions Ω_k and Ω_l are equal at the specified input boundary locations $\mathbf{x}_{k,l}$, as stated in the following theorem.

Theorem 1 *Suppose that $\Gamma_{k,l} \neq \emptyset$, i.e., Ω_k and Ω_l are neighboring. We have*

$$\begin{aligned}&\mathbb{E}[f_*^{(k)}|\boldsymbol{y}, \boldsymbol{\delta} = \mathbf{0}] = \mathbb{E}[f_*^{(l)}|\boldsymbol{y}, \boldsymbol{\delta} = \mathbf{0}] \text{ for } x_* \in \mathbf{x}_{k,l}, \text{ and} \\ &\mathbb{V}\mathrm{ar}[f_*^{(k)}|\boldsymbol{y}, \boldsymbol{\delta}] = \mathbb{V}\mathrm{ar}[f_*^{(l)}|\boldsymbol{y}, \boldsymbol{\delta}] \text{ for } x_* \in \mathbf{x}_{k,l}.\end{aligned} \quad (7)$$

Proof The difference of the two means

$$\begin{aligned}&\mathbb{E}[f_*^{(k)}|\boldsymbol{y},\boldsymbol{\delta}=\mathbf{0}]-\mathbb{E}[f_*^{(l)}|\boldsymbol{y},\boldsymbol{\delta}=\mathbf{0}]\\&=(\boldsymbol{c}_{*\mathcal{D}}^{(k)}-\boldsymbol{c}_{*\mathcal{D}}^{(l)}-(\boldsymbol{c}_{*\delta}^{(k)}-\boldsymbol{c}_{*\delta}^{(l)})\boldsymbol{C}_{\delta,\delta}^{-1}\boldsymbol{C}_{\mathcal{D}\delta}^{T})(\boldsymbol{C}_{\mathcal{D}\mathcal{D}}-\boldsymbol{C}_{\mathcal{D}\delta}\boldsymbol{C}_{\delta,\delta}^{-1}\boldsymbol{C}_{\mathcal{D}\delta}^{T})^{-1}\boldsymbol{y}.\end{aligned}\tag{8}$$

Let $\boldsymbol{V}$ denote the matrix with its ith row $\boldsymbol{c}_{*\mathcal{D}}^{(k)}-\boldsymbol{c}_{*\mathcal{D}}^{(l)}$ evaluated for x_* equal to the ith entry of $\mathbf{x}_{k,l}$, and let $\boldsymbol{W}$ denote the matrix with its ith row $\boldsymbol{c}_{*\delta}^{(k)}-\boldsymbol{c}_{*\delta}^{(l)}$ evaluated for x_* equal to the ith entry of $\mathbf{x}_{k,l}$. In addition, let $\boldsymbol{f}^{(k)}$ denote a column vector of the $f_*^{(k)}$ values for $x_*\in\mathbf{x}_{k,l}$ and $\boldsymbol{f}^{(l)}$ denote a column vector of the $f_*^{(l)}$ values for $x_*\in\mathbf{x}_{k,l}$. Using (8), we have

$$\mathbb{E}[\boldsymbol{f}^{(k)}-\boldsymbol{f}^{(l)}|\boldsymbol{y},\boldsymbol{\delta}]=(\boldsymbol{V}-\boldsymbol{W}\boldsymbol{C}_{\delta,\delta}^{-1}\boldsymbol{C}_{\mathcal{D}\delta}^{T})(\boldsymbol{C}_{\mathcal{D}\mathcal{D}}-\boldsymbol{C}_{\mathcal{D}\delta}\boldsymbol{C}_{\delta,\delta}^{-1}\boldsymbol{C}_{\mathcal{D}\delta}^{T})^{-1}\boldsymbol{y}.$$

Note that we can write $\boldsymbol{V}=\boldsymbol{\Theta}\boldsymbol{C}_{\mathcal{D}\delta}^{T}$ and $\boldsymbol{W}=\boldsymbol{\Theta}\boldsymbol{C}_{\delta,\delta}$ for some non-random matrix $\boldsymbol{\Theta}$. Therefore,

$$\begin{aligned}\mathbb{E}[\boldsymbol{f}^{(k)}-\boldsymbol{f}^{(l)}|\boldsymbol{y},\boldsymbol{\delta}]&=\boldsymbol{\Theta}(\boldsymbol{C}_{\mathcal{D}\delta}^{T}-\boldsymbol{C}_{\delta,\delta}\boldsymbol{C}_{\delta,\delta}^{-1}\boldsymbol{C}_{\mathcal{D}\delta}^{T})(\boldsymbol{C}_{\mathcal{D}\mathcal{D}}-\boldsymbol{C}_{\mathcal{D}\delta}\boldsymbol{C}_{\delta,\delta}^{-1}\boldsymbol{C}_{\mathcal{D}\delta}^{T})^{-1}\boldsymbol{y}\\&=\mathbf{0}.\end{aligned}$$

Let $\boldsymbol{Z}$ denote the matrix with its jth column $\boldsymbol{c}_{\mathcal{D}*}^{(k)}+\boldsymbol{c}_{\mathcal{D}*}^{(l)}$ evaluated for x_* equal to the jth entry of $\mathbf{x}_{k,l}$.

$$\begin{aligned}\mathbb{V}\mathrm{ar}[\boldsymbol{f}^{(k)}|\boldsymbol{y},\boldsymbol{\delta}]-\mathbb{V}\mathrm{ar}[\boldsymbol{f}^{(l)}|\boldsymbol{y},\boldsymbol{\delta}]&=\boldsymbol{\Theta}(\boldsymbol{C}_{\mathcal{D}\delta}^{T}-\boldsymbol{C}_{\delta,\delta}\boldsymbol{C}_{\delta,\delta}^{-1}\boldsymbol{C}_{\mathcal{D}\delta}^{T})(\boldsymbol{C}_{\mathcal{D}\mathcal{D}}-\boldsymbol{C}_{\mathcal{D}\delta}\boldsymbol{C}_{\delta,\delta}^{-1}\boldsymbol{C}_{\mathcal{D}\delta}^{T})^{-1}\boldsymbol{Z}\\&=\mathbf{0}.\end{aligned}$$

■

From the preceding discussion, our proposed approach enforces that the two local GP models for two neighboring local regions have the same posterior predictive means and variances at the chosen set of boundary points corresponding to the pseudo observations. We view this as patching together the independent local models in a nearly continuous way. The chosen sets of boundary points serve as the stitches when patching together the pieces, and the more boundary points are chosen, the closer the models are to being continuous over the entire boundary. In light of this, we refer to the approach as *patchwork kriging*.

3.1 Illustrative Example

To illustrate how patchwork kriging changes the model predictions (relative to a set of independent GP models over each region, with no continuity conditions imposed), we designed the following simple simulation study; we will present more comprehensive simulation comparisons and analyses in Sections 4 and 5. We generated a data set of 6,000 noisy observations

$$y_i=f(x_i)+\epsilon_i\quad\text{ for } i=1,\ldots,6000,$$

from a zero-mean Gaussian process with an exponential covariance function of scale one and variance 10, where $x_i\sim\text{Uniform}(0,10)$ and $\epsilon_i\sim\mathcal{N}(0,1)$ are independently sampled, and

$f(x_i)$ is simulated by the R package `RandomField`. Three hundred of the 6,000 observations were randomly selected as the training data $\mathcal{D}$, while the remaining 5,700 were reserved for test data. Figure 2 illustrates how the patchwork kriging predictor changes for different K, relative to the global GP predictor and the regular local GP predictor with no continuity conditions across regions. As the number of regions (K) increases, the regular local GP predictor deviates more from the global GP predictor. The test prediction mean square error (MSE) for the regular local GP predictor at the 5,700 test locations is 0.0137 for $K = 4$, 0.0269 for $K = 8$, 0.0594 for $K = 16$, and 0.1268 for $K = 32$. In comparison, patchwork kriging substantially improves the test MSE to 0.0072 for $K = 4$, 0.0123 for $K = 8$, 0.0141 for $K = 16$, and 0.0301 for $K = 32$.

3.2 Computational Expense and Sparsity

The computational expense of patchwork kriging is much less than that of the global GP regression. The computational expense of evaluating the predictive mean and variance for the patchwork kriging model is dominated by the inversion of $(\boldsymbol{C}_{\mathcal{D}\mathcal{D}} - \boldsymbol{C}_{\mathcal{D}\delta}\boldsymbol{C}_{\delta,\delta}^{-1}\boldsymbol{C}_{\mathcal{D}\delta}^T)$. Note that

$$(\boldsymbol{C}_{\mathcal{D}\mathcal{D}} - \boldsymbol{C}_{\mathcal{D}\delta}\boldsymbol{C}_{\delta,\delta}^{-1}\boldsymbol{C}_{\mathcal{D}\delta}^T)^{-1} = \boldsymbol{C}_{\mathcal{D}\mathcal{D}}^{-1} + \boldsymbol{C}_{\mathcal{D}\mathcal{D}}^{-1}\boldsymbol{C}_{\mathcal{D}\delta}(\boldsymbol{C}_{\delta,\delta} - \boldsymbol{C}_{\mathcal{D}\delta}^T\boldsymbol{C}_{\mathcal{D}\mathcal{D}}^{-1}\boldsymbol{C}_{\mathcal{D}\delta})^{-1}\boldsymbol{C}_{\mathcal{D}\delta}^T\boldsymbol{C}_{\mathcal{D}\mathcal{D}}^{-1}.$$

and that $\boldsymbol{C}_{\mathcal{D}\mathcal{D}}^{-1}$ is a block diagonal matrix with the kth block size equal to the size of $\mathcal{D}_k$. If the size of each $\mathcal{D}_k$ is M, evaluating $\boldsymbol{C}_{\mathcal{D}\mathcal{D}}^{-1}$ requires only inverting K matrices of size $M \times M$, and its expense is $O(KM^3)$. Given $\boldsymbol{C}_{\mathcal{D}\mathcal{D}}^{-1}$, evaluating $\boldsymbol{C}_{\mathcal{D}\mathcal{D}}^{-1}\boldsymbol{C}_{\mathcal{D}\delta}$ adds $O(KBM^2)$ to the computational expense. In addition, $\boldsymbol{C}_{\delta,\delta} - \boldsymbol{C}_{\mathcal{D}\delta}^T\boldsymbol{C}_{\mathcal{D}\mathcal{D}}^{-1}\boldsymbol{C}_{\mathcal{D}\delta}$ is a sparse matrix, and the number of non-zero elements in it is proportional to $O(B^2K)$, so its inversion adds $O(K^{1.5}B^3)$ to the expense. Therefore, the total computational expense of the patchwork kriging predictor is $O(KM^3 + KBM^2 + K^{1.5}B^3)$. Typically, $B \ll M$, in which case the complexity is $O(KM^3 + K^{1.5}B^3)$. Since $M \approx N/K$, the complexity can be written as $O(N^3/K^2 + K^{1.5}B^3)$. Given a fixed data size N, more splits of the regression domain will reduce the computation time due to the first term in the complexity, but too much increase would increase the second term, which will be shown later in Section 5.

4. Implementation Details

The performance of the proposed patchwork kriging method depends on the choice of tuning parameters, including the number of partitions (K) and the number (B) and locations ($\mathbf{x}_{k,l}$) of the pseudo observations. This section presents guidelines for these choices. Choosing the locations of pseudo observations is related to the choice of domain partitioning. In Section 4.1, we discuss the choices of the locations and partitioning together. In Section 4.2, we discuss the choice of K and B and illustrate with numerical experiments. Section 4.3 explains how to learn the parameters of a covariance function.

4.1 Choice of Ω_k and $\mathbf{x}_{k,l}$

There are many existing methods to partition a large set of data into smaller pieces. The simplest spatial partitioning is a uniform grid partitioning that divides a domain into uniform grids and splits data accordingly (Park et al., 2011; Park and Huang, 2016). This

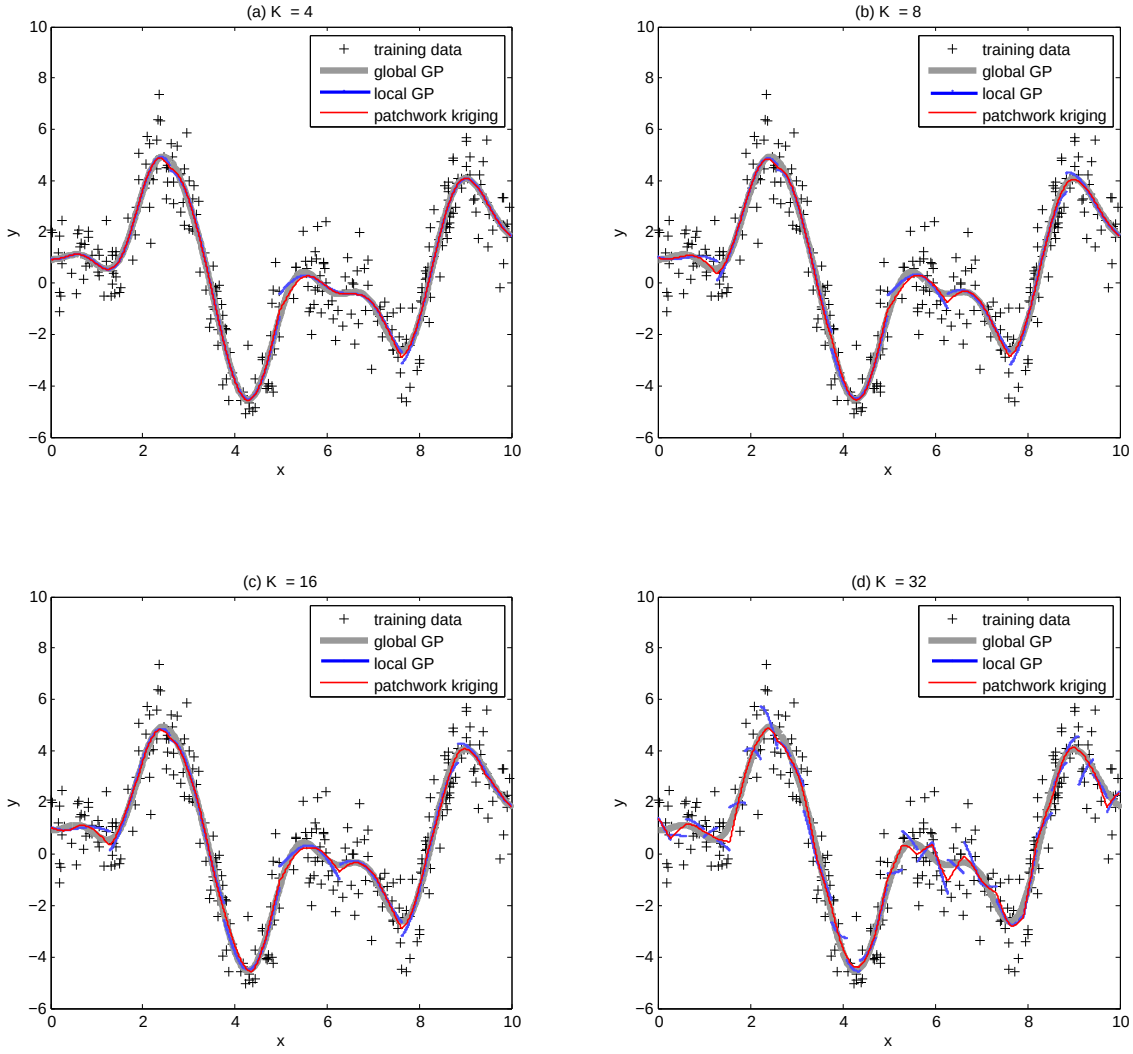

Figure 2: Example illustrating the patchwork kriging predictor, together with the global GP predictor and the regular local GP predictor with no continuity constraints. K is the number of local regions.

is simple and effective if the data are uniformly distributed over a low dimensional space. However, if the input dimension is high, it would either generate too many regions or it would produce many sparse regions that contain very few or no observations, and the latter also happens when the data are non-uniformly distributed; see examples in Figure 3-(c) and Figure 3-(e). Shen et al. (2006) used a kd-tree for spatial partitioning of unevenly distributed data points in a high dimensional space. A kd-tree is a recursive partitioning scheme that recursively bisects the subspaces along one chosen data dimension at a time. Later, McFee and Lanckriet (2011) generalized it to the spatial tree. Starting with a level

0 space $\Omega_1^{(0)}$ equal to the entire regression domain, the spatial tree recursively bisects each of level s spaces into two level $s+1$ spaces. Let $\Omega_j^{(s)} \in \mathbb{R}^d$ denote the jth region in the level s space. It is bisected into two level $s+1$ spaces as

$$\Omega_{2j-1}^{(s+1)} = \{\boldsymbol{x} \in \Omega_j^{(s)} : \boldsymbol{v}_{j,s}^T \boldsymbol{x} \leq \nu\} \text{ and } \Omega_{2j}^{(s+1)} = \{\boldsymbol{x} \in \Omega_j^{(s)} : \boldsymbol{v}_{j,s}^T \boldsymbol{x} > \nu\}. \tag{9}$$

Each of $\Omega_{2j-1}^{(s+1)}$ and $\Omega_{2j}^{(s+1)}$ will be further partitioned in the next level using the same procedure. The choice of the linear projection vector $\boldsymbol{v}_{j,s}$ depends on the distribution of the local data belonging to $\Omega_j^{(s)}$. For example, it can be the first principal component direction of the local data. The value for ν is chosen so that $\Omega_{2j-1}^{(s+1)}$ and $\Omega_{2j}^{(s+1)}$ have an equal number of observations. In this sense, the subregions at the same level are equally sized, which helps to level off the computation times of the local models. When the spatial tree is applied on data uniformly distributed over a rectangular domain, it produces a uniform grid partitioning; see examples in Figure 3-(b). The spatial tree is more effective than the grid partitioning when data is unevenly distributed; see examples in Figure 3-(d) and Figure 3-(f).

In this work, we use a spatial tree with the principal component (PC) direction for $\boldsymbol{v}_{j,s}$. Bisecting a space along the PC direction has effects of minimizing the area of the interfacial boundaries in between the two bisected regions, so the number of the pseudo observations necessary for connecting the two regions can be minimized. The maximum level of the recursive partitioning depends on the choice of K via $s_{max} = \lfloor \log_2 K \rfloor$. The choices of K and B will be discussed in the next section. Given B, the pseudo observations $\mathbf{x}_{k,l}$ are randomly generated from an uniform distribution over the intersection of the hyper-plane $\boldsymbol{v}_{j,s}^T \boldsymbol{x} = \nu$ and the level s region $\Omega_j^{(s)}$.

4.2 Choice of K and B

This section presents a simulation study to show how the patchwork kriging works with different choices of K and B, which helps to provide practical guidelines for their choices. We generated a synthetic data set in 2-d using the R package `RandomField`, and we denote this dataset by `synthetic-2d`. `synthetic-2d` consists of 8,000 noisy observations from a zero-mean Gaussian process with an exponential covariance function of scale one and variance 10,

$$y_i = f(x_i) + \epsilon_i \quad \text{for } i = 1, \ldots, 8000,$$

where $x_i \sim \text{Uniform}([0,6] \times [0,6])$ and $\epsilon_i \sim \mathcal{N}(0,1)$ were independently sampled. For $N < 10,000$ roughly, we can compute the predictive mean and variance of the optimal GP predictor (i.e., using the true exponential GP model globally). We used these optimal prediction values as a benchmark against which to judge the accuracy of the patchwork kriging predictions.

We evaluated the ground truth means and variances at 800 test locations and compared the values with the predicted values of the patchwork kriging with different choices of K and B. The first 500 test locations are uniformly distributed over the entire regression domain, and the remaining 300 test locations are selected from the interfacial boundaries in between neighboring local regions. The mean square differences (MSD) between the benchmark predictions and the patchwork kriging predictions are computed for each of the

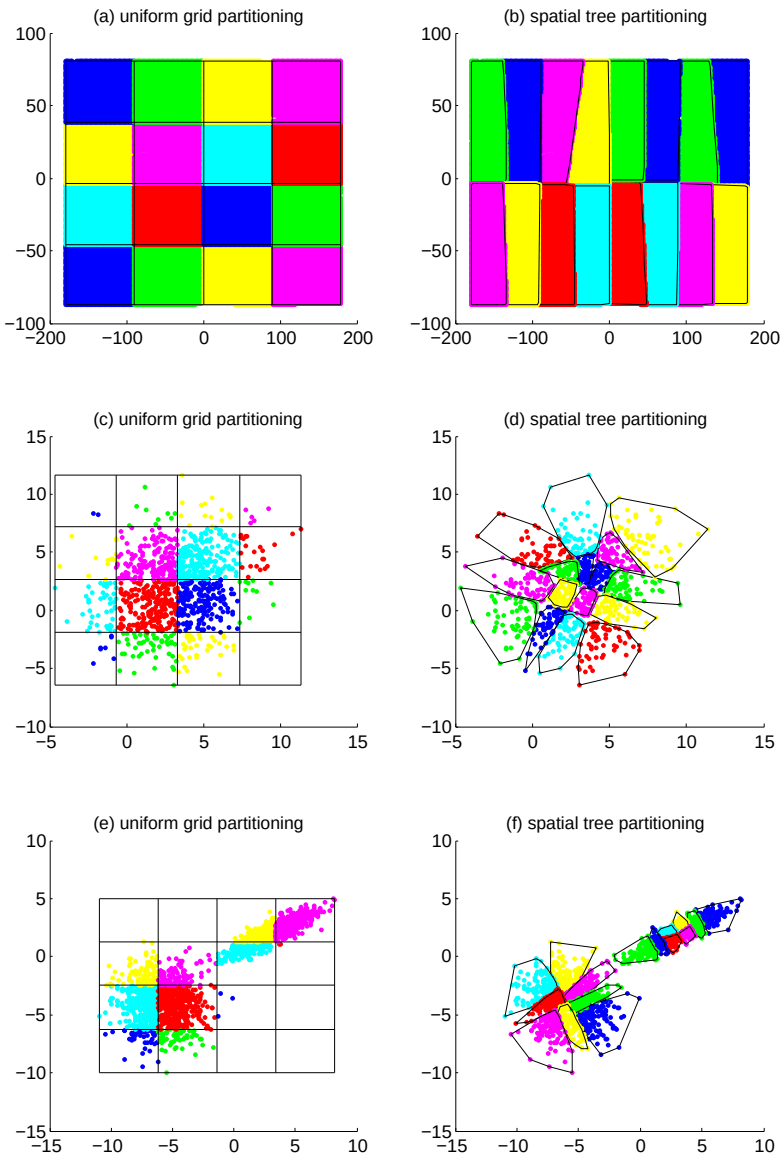
(a) uniform grid partitioning
(b) spatial tree partitioning
(c) uniform grid partitioning
(d) spatial tree partitioning
(e) uniform grid partitioning
(f) spatial tree partitioning

Figure 3: Comparison of two spatial partitioning schemes: a uniform grid (left panel) and a spatial tree (right panel).

predictive mean and variance predictions. Figure 4-(a) through Figure 4-(d) summarize the results. When B is fixed, the MSDs for the predictive mean and variance decrease as K decreases. For any K, increasing B significantly decreases the MSD, while the increase of B beyond five provides little further accuracy gains. Therefore, with B above a certain value, K should not be a deciding factor for the prediction accuracy.

On the other hand, from Figure 4-(e), we can see that the deciding factor for computation time is K rather than B. In general, larger K results in shorter computation times, while too large K increases computation time. This is expected from the analysis of computational complexity in Section 3. The accuracy and time analysis suggests choosing K that gives the smallest computation time, while fixing B above 5 to maintain a good prediction accuracy. According to Section 3, the actual computation time would be proportional to $N^3/K^2 + qK^{1.5}B^3$ for an unknown positive constant q. Taking its derivative with respect to K, we have the first order necessary condition for K that minimizes the time,

$$-2N^3/K^3 + q1.5K^{0.5}B^3 = 0 \quad \Rightarrow \quad K = \left(\frac{N}{qB}\right)^{3/3.5}.$$

Our many numerical experiments suggest approximately $B = 7$ and $q = 30$, which gives the choice $K = (N/210)^{3/3.5}$. For example, applying this suggestion rule to the three spatial datasets used in Section 5 gave the shortest computation times or the second shortest computation times (and very close to the shortest) among all configurations of K and B which we have tried.

We used the same `synthetic-2d` data to illustrate how two local GP models for two neighboring local regions change as B changes. We first partitioned the dataset into 128 local regions as shown in Figure 5. For evaluation purposes, we considered test points that fell on the boundary cutting the entire regression domain into two (indicated by the black solid line in Figure 5), and sampled 201 test points uniformly over this boundary. For each point, we get two mean predictions from the two local patchwork kriging models that straddle the boundary at that point. We compared the two mean predictions to each other for different choices of B and also compared them with the optimal global GP prediction, i.e., the prediction using the true GP covariance function and the entire dataset globally without spatial partitioning. Figure 6 shows the comparison. When $B = 0$, the two local models exhibits significant differences in their mean predictions. The differences decrease as B increases, and become negligible when $B \geq 5$. The mean predictions are also very close to the exact GP predictions.

4.3 Hyper-parameter Learning

The hyperparameters of the covariance function $c_k(\cdot, \cdot)$ determines the correlation among the values of $f(x)$, which has significant effect on the accuracy of a Gaussian process regression. We estimate all correlation parameters using maximum likelihood estimation (MLE) by minimizing the negative log-likelihood

$$NL(\theta) = \frac{N}{2}\log(2\pi) + \frac{1}{2}\log\left|\begin{array}{cc} \boldsymbol{C}_{\mathcal{D}\mathcal{D}} & \boldsymbol{C}_{\mathcal{D},\delta} \\ \boldsymbol{C}_{\delta,\mathcal{D}} & \boldsymbol{C}_{\delta,\delta} \end{array}\right| + \frac{1}{2}[\boldsymbol{y}^T \mathbf{0}^T] \left[\begin{array}{cc} \boldsymbol{C}_{\mathcal{D}\mathcal{D}} & \boldsymbol{C}_{\mathcal{D},\delta} \\ \boldsymbol{C}_{\delta,\mathcal{D}} & \boldsymbol{C}_{\delta,\delta} \end{array}\right]^{-1} \left[\begin{array}{c} \boldsymbol{y} \\ \mathbf{0} \end{array}\right] \quad (10)$$

Note that we have augmented the data to included the pseudo observations $\boldsymbol{\delta} = \mathbf{0}$ in the likelihood expression, which results in a better behaved likelihood by imposing some continuity across regions. Using partitioned matrix results similar to the matrix inversion (5), the determinant expression becomes

$$\left|\begin{array}{cc} \boldsymbol{C}_{\mathcal{D}\mathcal{D}} & \boldsymbol{C}_{\mathcal{D},\delta} \\ \boldsymbol{C}_{\delta,\mathcal{D}} & \boldsymbol{C}_{\delta,\delta} \end{array}\right| = |\boldsymbol{C}_{\mathcal{D}\mathcal{D}}||\boldsymbol{C}_{\delta,\delta} - \boldsymbol{C}_{\delta,\mathcal{D}}\boldsymbol{C}_{\mathcal{D}\mathcal{D}}^{-1}\boldsymbol{C}_{\mathcal{D},\delta}|,$$

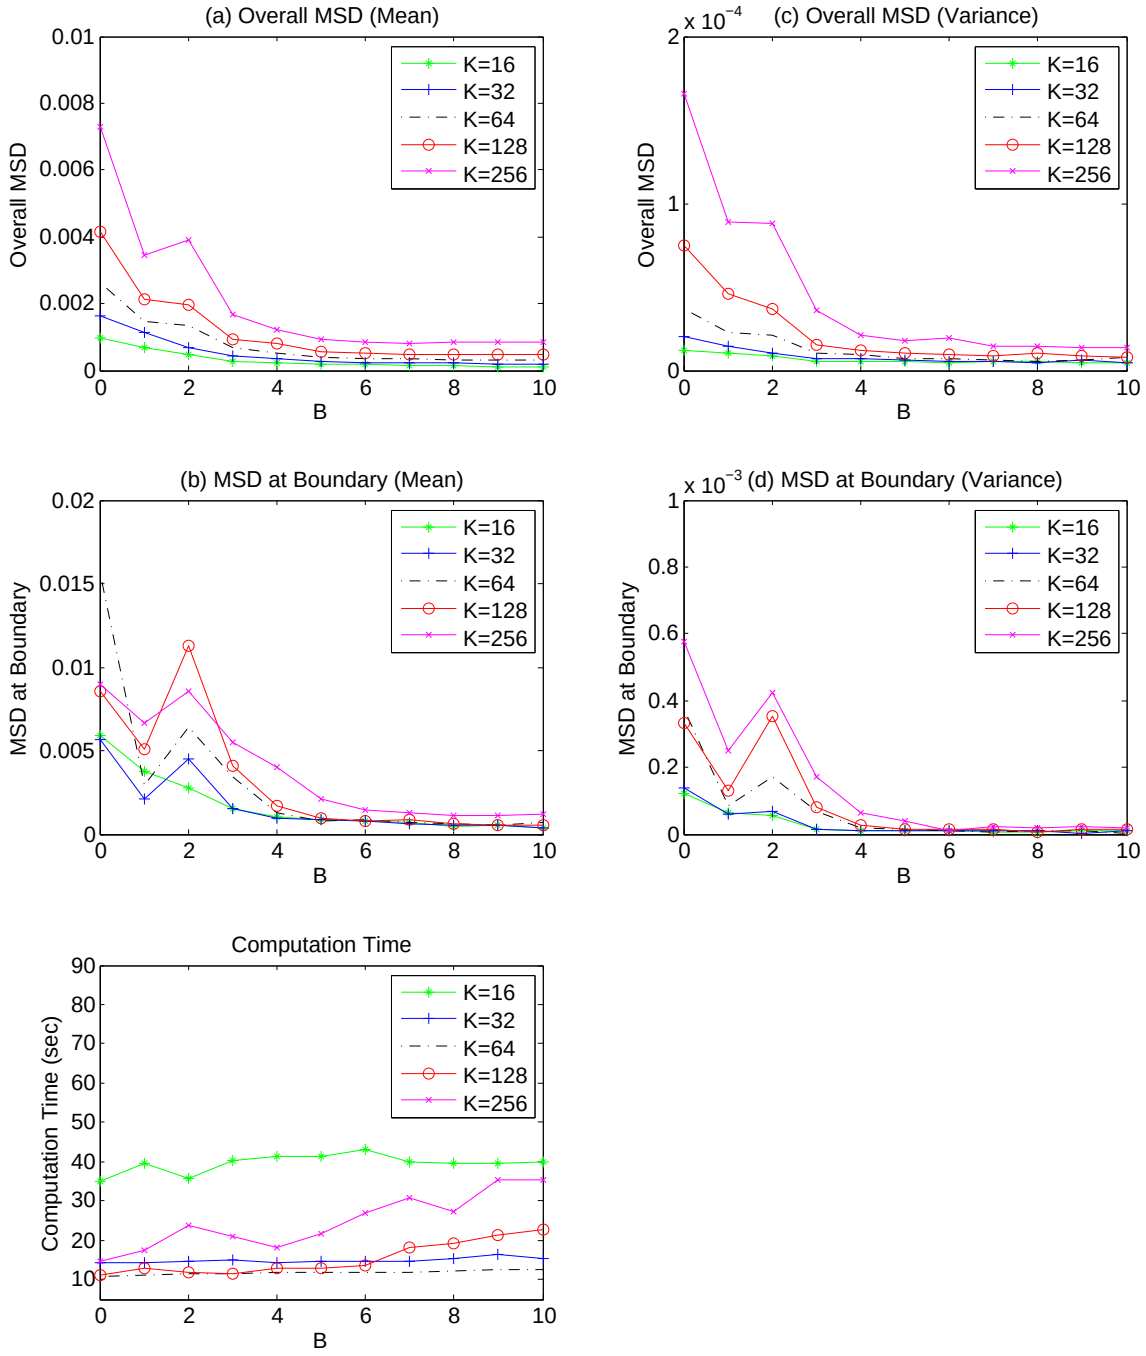

Figure 4: Accuracy of the predictive mean and variance estimates versus B and K. In plots (a)-(d), the MSE measures the mean square difference between the values from the optimal global GP predictor and the values from the patchwork kriging predictor.

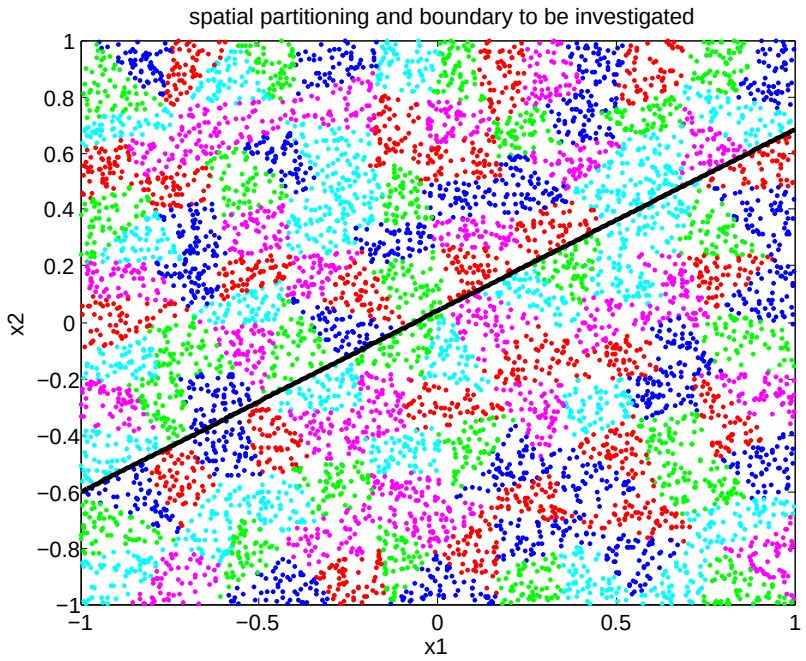

Figure 5: The entire regression domain was spatially partitioned into 128 local regions, which are distinguished by their colors. The black solid line cutting through the entire space is the interfacial boundary that we selected to study the behavior of the patchwork kriging at interfacial boundaries.

and the quadratic term becomes

$$[\boldsymbol{y}^T \boldsymbol{0}^T] \begin{bmatrix} \boldsymbol{C}_{\mathcal{D}\mathcal{D}} & \boldsymbol{C}_{\mathcal{D},\delta} \\ \boldsymbol{C}_{\delta,\mathcal{D}} & \boldsymbol{C}_{\delta,\delta} \end{bmatrix}^{-1} \begin{bmatrix} \boldsymbol{y} \\ \boldsymbol{0} \end{bmatrix} = \boldsymbol{y}^T (\boldsymbol{C}_{\mathcal{D}\mathcal{D}} - \boldsymbol{C}_{\mathcal{D}\delta} \boldsymbol{C}_{\delta,\delta}^{-1} \boldsymbol{C}_{\delta\mathcal{D}})^{-1} \boldsymbol{y},$$

which further simplifies to

$$\boldsymbol{y}^T \boldsymbol{C}_{\mathcal{D}\mathcal{D}}^{-1} \boldsymbol{y} + \boldsymbol{y}^T \boldsymbol{C}_{\mathcal{D}\mathcal{D}}^{-1} \boldsymbol{C}_{\mathcal{D}\delta} (\boldsymbol{C}_{\delta,\delta} - \boldsymbol{C}_{\delta,\mathcal{D}} \boldsymbol{C}_{\mathcal{D}\mathcal{D}}^{-1} \boldsymbol{C}_{\mathcal{D},\delta})^{-1} \boldsymbol{C}_{\delta\mathcal{D}} \boldsymbol{C}_{\mathcal{D}\mathcal{D}}^{-1} \boldsymbol{y}.$$

Thus, the same computational tricks used when calculating the patchwork kriging predictions can be used to reduce the computational expense when estimating the covariance parameters via MLE.

5. Numerical Examples

In this section, we use five real datasets to evaluate the patchwork kriging and compare it with the state-of-the-art, including (Park and Huang, 2016, PGP), a Gaussian Markov random field approximation (Lindgren et al., 2011, GMRF), a robust Bayesian committee machine (Deisenroth and Ng, 2015, RBCM), and a partial information conditional approach (Snelson and Ghahramani, 2007, PIC).

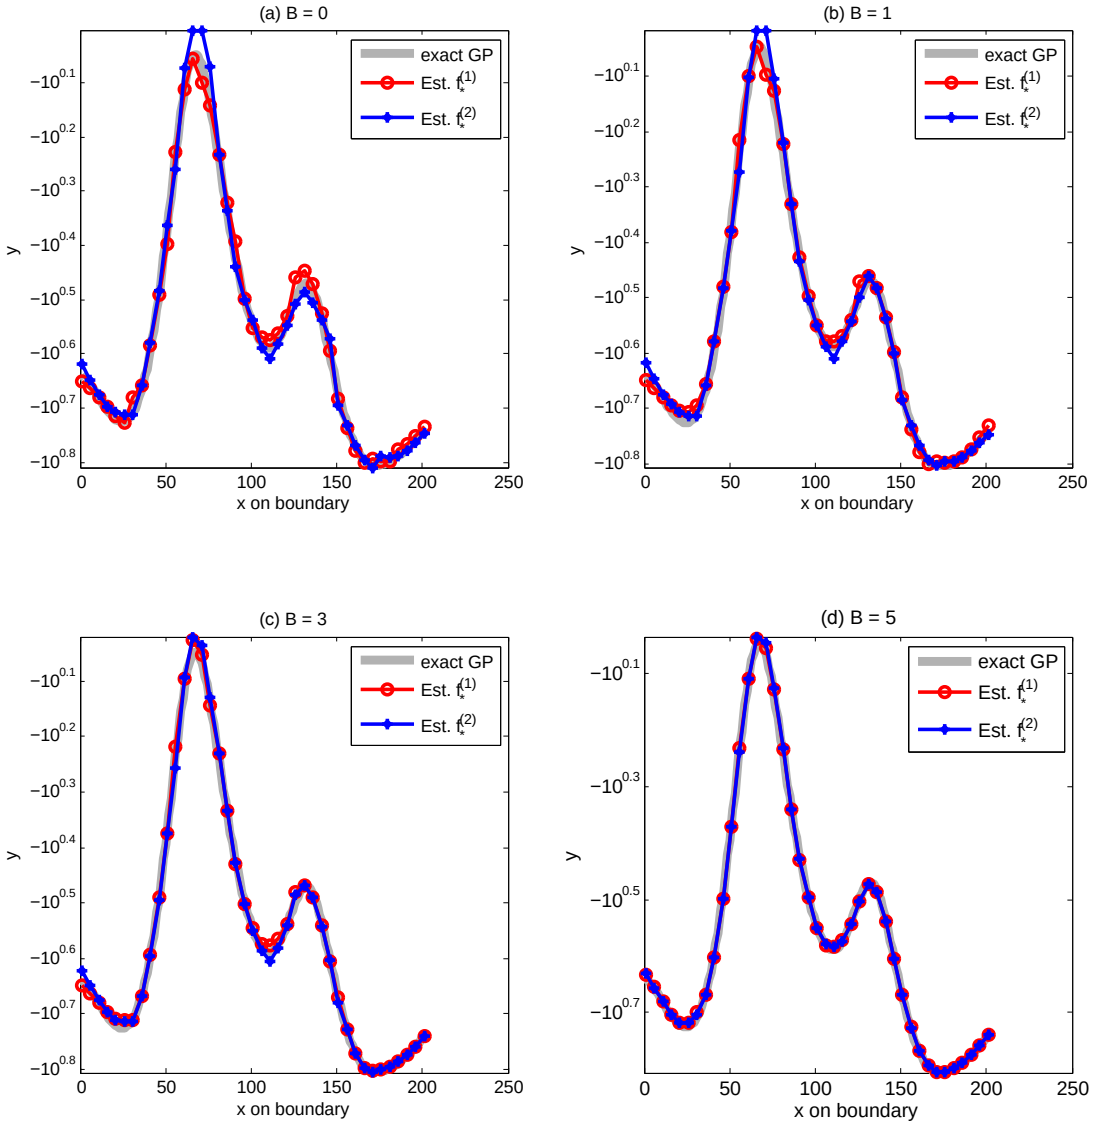

Figure 6: Comparison of the patchwork kriging mean predictions for the two local models in the two neighboring regions straddling the interfacial boundary. The patchwork kriging predictors are using B = 0, 1, 3, and 5 for panels (a), (b), (c), and (d), respectively. The horizontal axes represent the location along the solid interfacial boundary line shown in Figure 5. $f_*^{(1)}$ and $f_*^{(2)}$ denote the mean predictions of the two local models on each side of the solid boundary line. As B increases, the two local predictions converge to each other, and the converged values are very close to the benchmark predictor achieved using the true GP model globally.

5.1 Data Sets and Evaluation Criteria

We considered five real data sets: two spatial datasets in 2-d with different spatial distributions of observations, one additional spatial dataset with a very large data size, and two high dimensional datasets, one with 9 dimensions and the other with 21 dimensions.

The first data set, `TCO`, contains 48,331 observations collected by the NIMBUS-7/TOMS satellite, which measures the total column of ozone over the globe on Oct 1 1988. The second spatial data set, `TCO.L2`, contains the same data but at many more locations (182,591 locations). The third data set, `ICETHICK`, contains ice thickness measurements at 32,481 locations on the western Antarctic ice sheet and is available at `http://nsidc.org/`. For the first two spatial datasets, the observations are uniformly spread over the input dimensions, while the third dataset has many sparse regions with very few observations.

The first high dimensional dataset, `PROTEIN`, has nine input variables that describe the tertiary structure of proteins and one response variable that describes the physiochemical properties of protein. These data, which are available at `https://archive.ics.uci.edu/ml/datasets`, consist of 45,730 observations. Like typical high dimensional datasets, the measurements are embedded on a low dimensional subspace of the entire domain.

The second high dimensional dataset, `SARCOS`, has 21 input variables that describe the positions, velocities and accelerations of seven robot arm locations, and there are seven separate response variables. We only use the first response variable for this numerical study. The dataset, which is available at `http://www.gaussianprocess.org/gpml/data/`, contains 44,484 training observations and 4,449 test observations.

Using the five data sets, we compare the computation time and prediction accuracy of patchwork kriging with other methods. We randomly split each data set into a training set containing 90% of the total observations and a test set containing the remaining 10% of the observations. To compare the computational efficiency of methods, we measure total computation times. For comparison of prediction accuracy, we measure two performance metrics on the test data, denoted by $\{(x_t, y_t) : t = 1, \ldots, T\}$, where T is the size of the test set. Let μ_t and σ_t^2 denote the estimated posterior predictive mean and variance at location x_t. The first measure is the mean squared error (MSE)

$$\text{MSE} = \frac{1}{T}\sum_{t=1}^{T}(y_t - \mu_t)^2, \tag{11}$$

which measures the accuracy of the mean prediction μ_t at location x_t. The second measure is the negative log predictive density (NLPD)

$$\text{NLPD} = \frac{1}{T}\sum_{t=1}^{T}\left[\frac{(y_t - \mu_t)^2}{2\sigma_t^2} + \frac{1}{2}\log(2\pi\sigma_t^2)\right]. \tag{12}$$

The NLPD quantifies the degree of fitness of the estimated predictive distribution $\mathcal{N}(\mu_t, \sigma_t^2)$ for the test data. A good fit can be achieved when both the predictive mean μ_t and the predictive variance σ_t^2 are accurate. These two criteria are used broadly in the GP regression literature. A smaller value of MSE or NLPD indicates better performance. All numerical experiments were performed on a desktop computer with Intel Xeon Processor W3520 and 6GB RAM.

The comparison was made for five different methods including patchwork kriging (our method), the method of (Park and Huang, 2016, PGP), the Gaussian Markov random field approximation (Lindgren et al., 2011, GMRF), robust Bayesian committee machine (Deisenroth and Ng, 2015, RBCM) and partial information conditional (Snelson and Ghahramani, 2007, PIC). Note that the PGP and the GMRF approaches cannot be applied for more than two input dimensions, and so were only compared for the three spatial datasets. We tried two covariance functions, a squared exponential covariance function and an exponential covariance function. Note that the PIC method does not work with an exponential covariance function because learning the pseudo inputs for the PIC method requires the derivative of a covariance function but an exponential covariance function is not differentiable. On the other hand, when an squared exponential covariance function is applied to the GMRF, the precision matrix construction is not straightforward. Therefore, we used a squared exponential covariance function for comparing the proposed approach with the PIC, RBCM, and PGP, while using an exponential covariance function for comparing it with the GMRF. For both of the cases, we used the entire training data set to estimate the hyperparameters.

5.2 Example 1: `TCO` Dataset

This data set has two input dimensions, and the inputs of the data are densely distributed over a rectangular domain. For patchwork kriging, we varied $B \in \{5, 7\}$ and $K \in \{64, 128, 256, 512\}$. The prediction accuracy of the PGP did not depend on the number of local regions K, so we fixed $K = 145$, while the number of finite element meshes per local region was varied from 5 to 40 with step size 5. For RBCM, we varied the number of local experts $K \in \{100, 150, 200, 250, 300, 600\}$. For PIC, K was varied over $\{100, 150, 200, 250, 300, 400\}$, and the number of pseudo inputs was also varied over $\{30, 50, 70, 80, 100, 150, 200, 250, 300\}$. The GMRF has no tuning parameter.

Figure 7 summarizes the performance results. The two panels in the top row compare the proposed approach to the PGP, RBCM and PIC approaches when a square exponential covariance function is used. The MSE plot shows the mean square error of the mean predictions for the test data. The proposed approach and the PGP approach had better MSE than the RBCM and PIC approaches, and the proposed approach also had better MSE than the PGP approach at the 100 second computation time. The NLPD plot shows the degree of fitness of the estimated predictive distribution to the test data, which is affected by both of the mean prediction and the variance prediction. The proposed approach uniformly outperforms the other methods, including the PGP. The two panels in the bottom row compare the proposed approach to the PGP and the GMRF approaches when an exponential covariance function is used. The PGP and the proposed approaches had comparable MSEs at larger computation times, whereas the proposed approach had much smaller MSE at the lower computation times, and the proposed approach outperformed the PGP and the GMRF in terms of NLPD. The GMRF had longer computation time than the PGP and the proposed method.

5.3 Example 2: `TCO.L2` Dataset

This dataset is about four times larger than the other datasets, so it was used to see how the methods work for larger datasets. For patchwork kriging, we varied $B \in \{3, 5\}$ and $K \in$

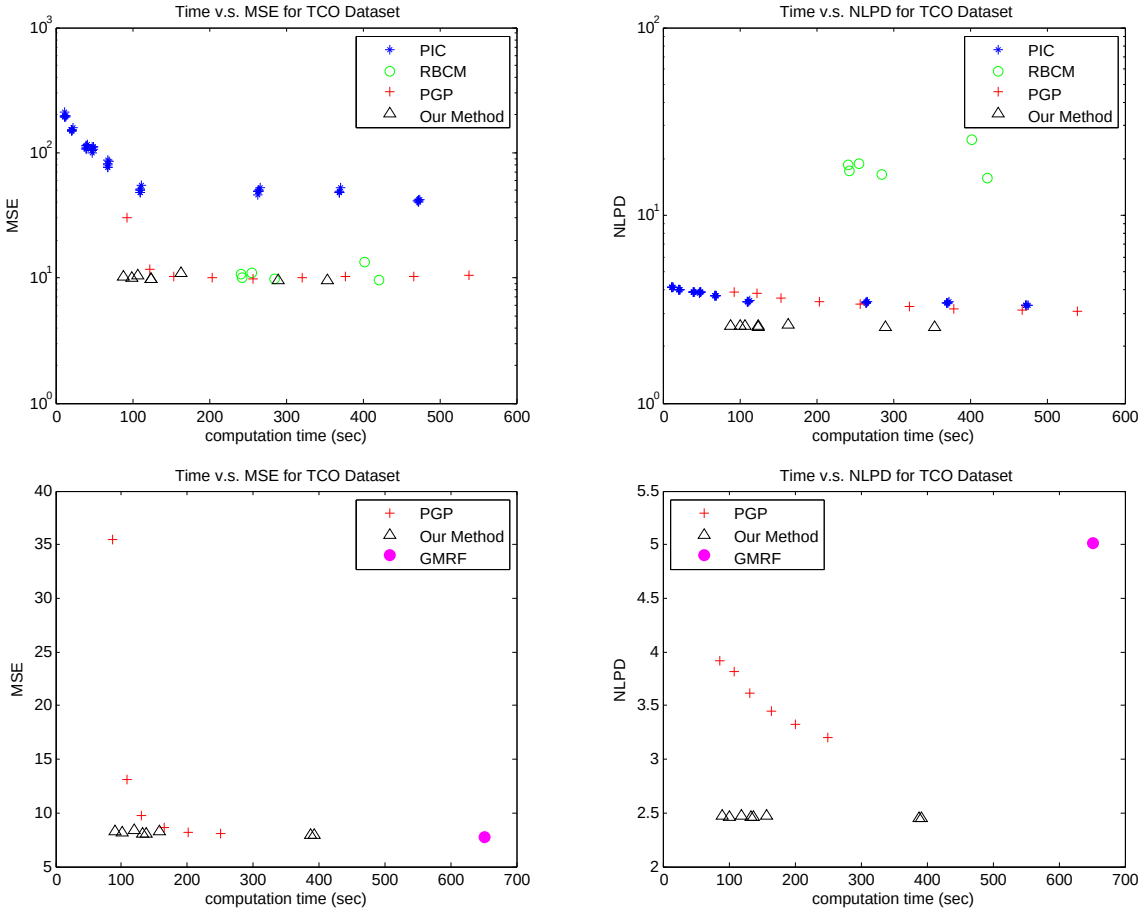

Figure 7: Prediction accuracy versus total computation time for the `TCO` data. Top panel: a squared exponential covariance function was used. Bottom panel: an exponential covariance function was used.

$\{256, 512, 1024\}$. The prediction accuracy of the PGP did not depend on the number of local regions K, so we fixed $K = 623$, while the number of finite element meshes per local region was varied from 5 to 25 with step size 5. For RBCM, we varied the number of local experts $K \in \{100, 150, 200, 250, 300, 600\}$. For PIC, K was varied over $\{100, 200, 300, 400, 600\}$, and the number of pseudo inputs was also varied over $\{50, 70, 80, 100, 150, 200, 300\}$.

Figure 8 shows the main results. The shortest computation time of RBCM (2319 seconds) was much longer than the longest time of the other compared methods, while its MSE was not competitive as well. Therefore we did not plot its results in the figure. For both of the square exponential and the exponential covariance functions, our approach and the PGP approach had comparable MSE. However, our approach significantly outperformed the PGP and PIC approaches in terms of the NLPD. This implies that our approach provides more accurate variance estimations.

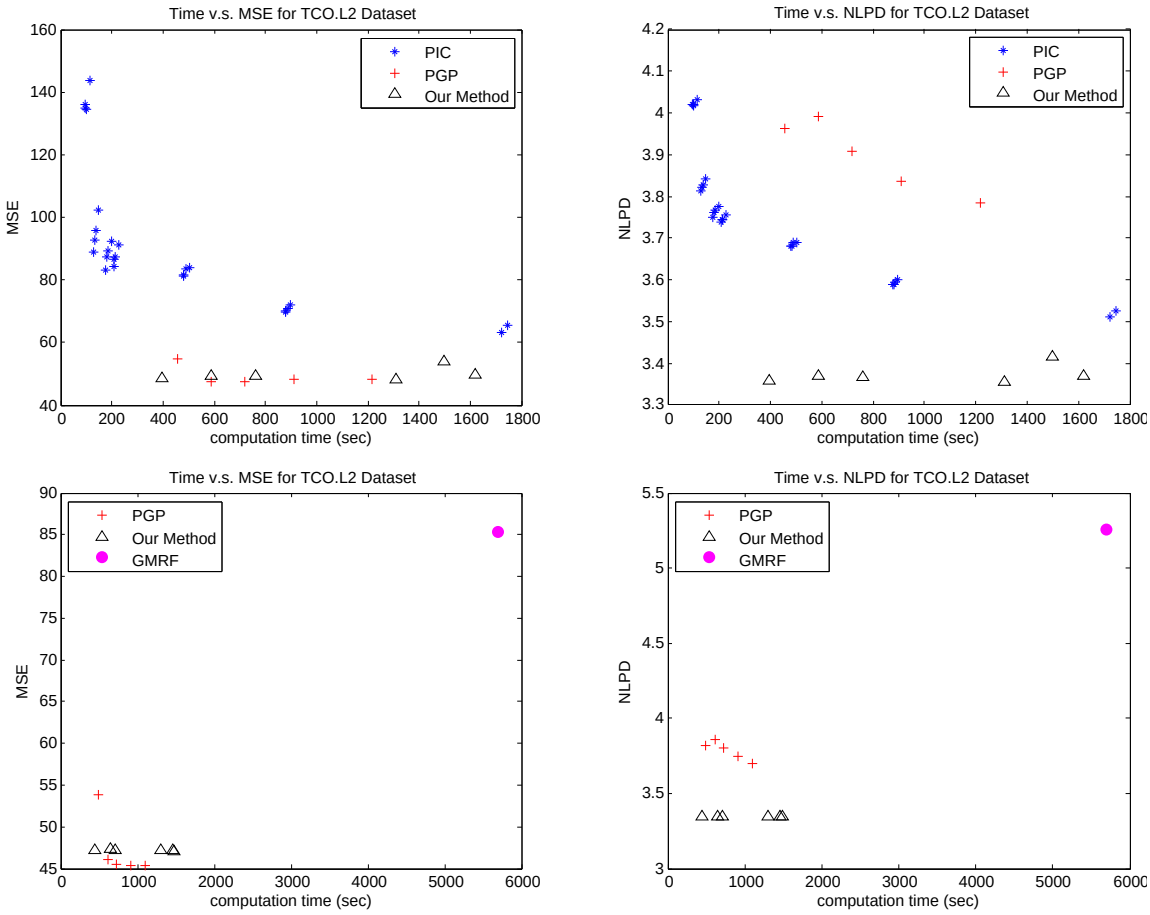

Figure 8: Prediction accuracy versus total computation time for the `TCO.L2` data. Top panel: a squared exponential covariance function was used. Bottom panel: an exponential covariance function was used.

5.4 Example 3: `ICETHICK` Dataset

One characteristic of this dataset is the presence of many spatial voids where there are no or very little data points. For patchwork kriging, we varied $B \in \{3, 5, 7\}$ and $K \in \{64, 128, 256, 512, 1024\}$. For the PGP, we used $K = 47$, while the number of finite element meshes per local region was varied from 5 to 40 with step size 5. For RBCM, we varied the number of local experts $K \in \{50, 100, 150, 200, 250, 300\}$. For PIC, M was varied over $\{50, 100, 150, 200\}$, and the number of pseudo inputs was also varied over $\{50, 100, 150, 200, 300, 400, 500, 600, 700\}$.

Figure 9 compares the MSE and NLPD performance of the methods. Again, the PGP approach and the proposed approach outperformed the other methods, and the proposed approach achieved the best accuracy with much less computation time than the PGP approach. In addition, the proposed approach uniformly outperformed the other methods in

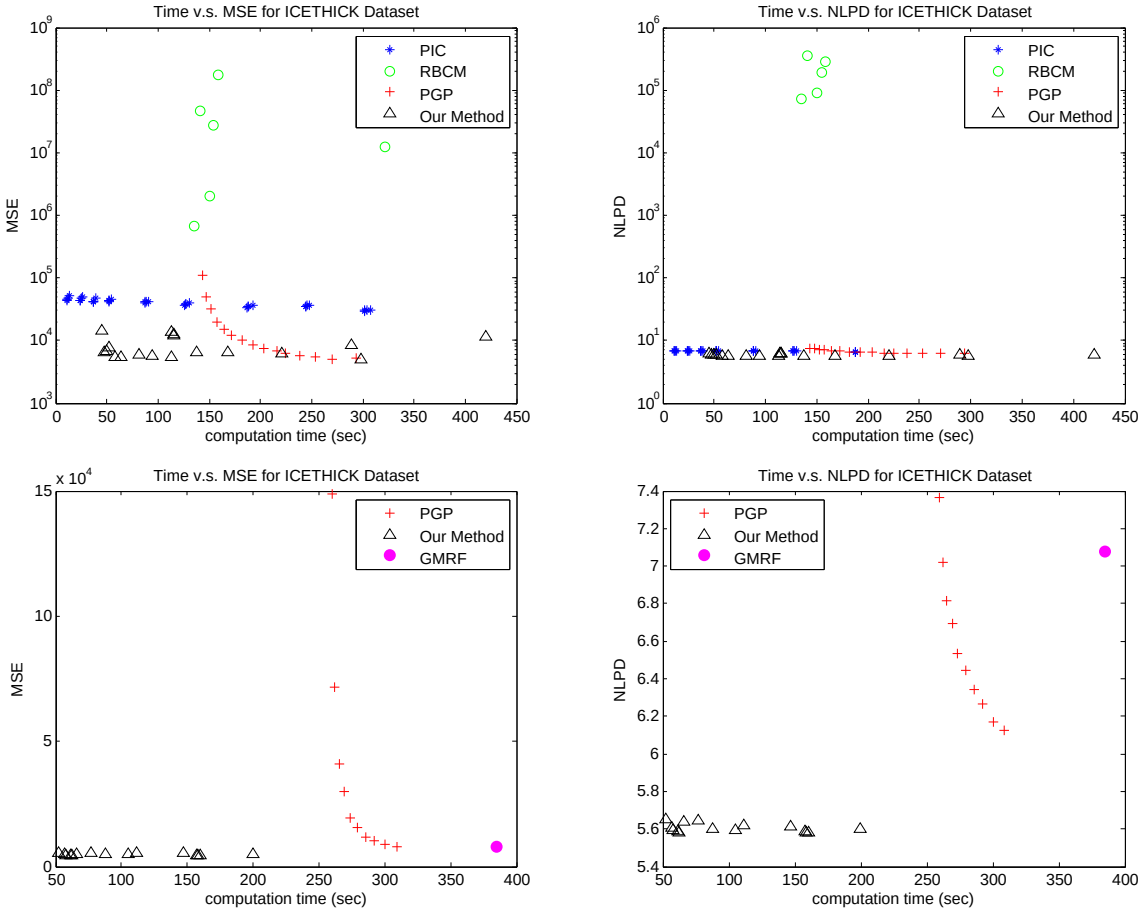

Figure 9: Prediction accuracy versus total computation time for the `ICETHICK` data. A squared exponential covariance function was used for the results in the top panel, while an exponential covariance function was used for the results in the bottom panel.

terms of the NLPD. In other words, the proposed approach gives a predictive distribution that better fits the test data.

5.5 Example 4: `PROTEIN` Dataset

Different from the previous datasets, this dataset features nine input variables. We will use this dataset to see how the proposed approach works for high dimensional cases. For the patchwork kriging, we varied $B \in \{2, 3, 4\}$ and varied $K \in \{64, 128, 256\}$; we have not included the results for larger B because a larger B increased the computation times of the patchwork kriging to a range incomparable to those of the other algorithms. The PGP and GMRF approaches do not work with high dimensional inputs, and so were not included in this comparison. For RBCM, we varied the number of local experts

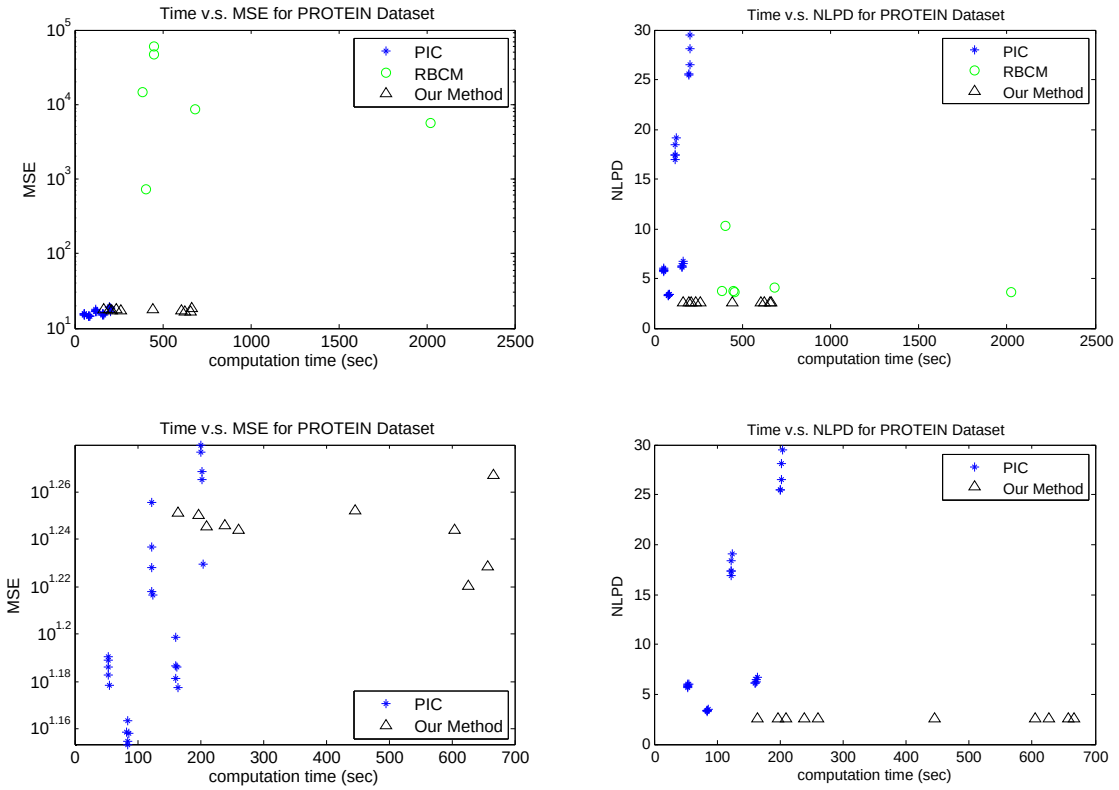

Figure 10: Prediction accuracy versus total computation time for the `PROTEIN` data. The upper panel compares all three methods. Since the performance of the PIC and our method was very close, the bottom panel provides a closer look of the PIC and our method.

$K \in \{100, 150, 200, 250, 300\}$. For PIC, K was varied over $\{100, 150, 200, 250, 300\}$, and the number of pseudo inputs (M) was also varied over $\{100, 150, 200, 250, 300\}$. In this comparison, we used a squared exponential covariance function for all three methods.

Figure 10 shows the main results. For this high dimensional dataset, the PIC approach outperformed our method in terms of the MSE performance, providing more accurate mean predictions. On the other hand, our method provided better NLPD performance, which implies that the predictive distribution estimated by our method was better fit to test data than that of the PIC. Figure 11 compares the predictive distributions estimated by the two methods. In the figure, the predicted mean ± 1.5 predicted standard deviations was plotted for 100 randomly chosen test observations. The interval for the PIC was overly narrow and excluded many of the test response observations, while the interval for our method more appropriately covered the majority of data, which is reflected in the better NLPD performance for our method. The percentages of the 4,573 test observations falling within the intervals was 50.47% for the PIC and 86.53% for our method. Note that the probability

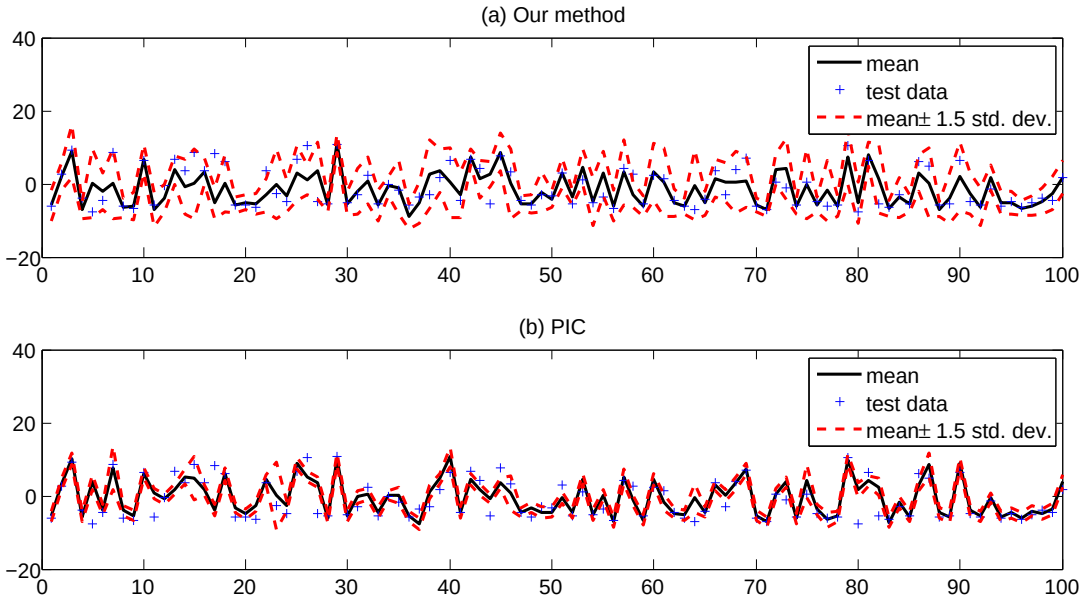

Figure 11: Comparison of the predictive distributions estimated by our method and by the PIC method for the PROTEIN data.

of a standard normal random variable within $\pm 1.5\sigma$ is 86.64%. Clearly, our method provides a better fit to the test data.

5.6 Example 5: SARCOS Dataset

This dataset has 21 input variables. For patchwork kriging, we fixed $B = 3$ since a further increase in B did not improve the performance, and we varied $K \in \{128, 256, 512, 1024\}$. Again, the PGP and GMRF approaches do not work with high dimensional inputs, and so were not included in this comparison. For the RBCM approach, we varied the number of local experts $K \in \{100, 150, 200, 250, 300\}$. For PIC, K was varied over $\{100, 150, 200, 250, 300\}$, and the number of pseudo inputs (M) was also varied over $\{100, 150, 200, 250, 300\}$. In this comparison, we used a squared exponential covariance function for all three methods.

Figure 12 summarizes the comparison of the MSEs and the NLPDs for the three methods. The MSE performances were comparable for all of the methods, while our approach provided a better fit to test data, which was evidenced by the smaller NLPD values of our approach. The PIC produced negative predictive variances for this dataset, so its NLPD values could not be calculated. In theory, the PIC approach should provide non-negative predictive variances with infinite precision. It evidently produced negative variances because of numerical errors. Our patchwork kriging approach did not experience any such numerical issues in any of the examples, and it appears to be more numerically stable than the PIC approach.

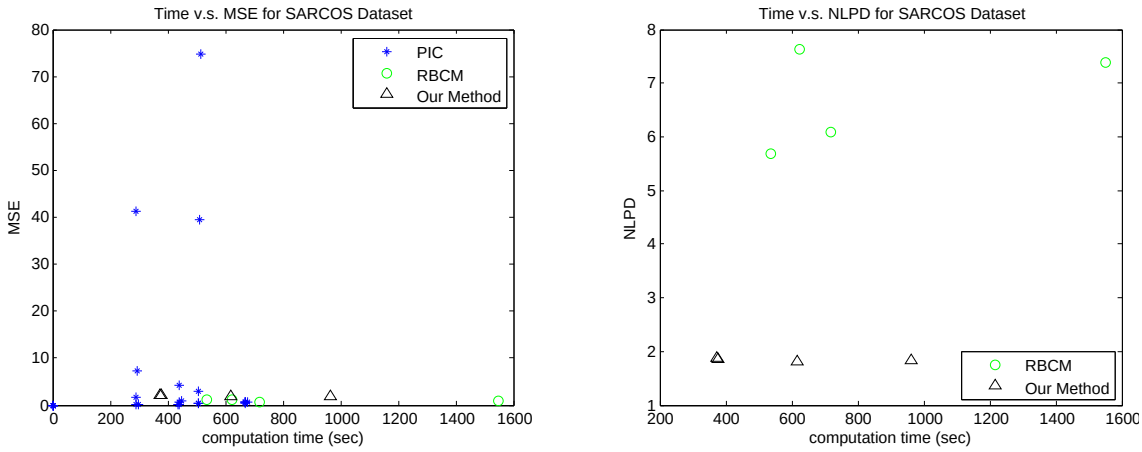

Figure 12: Prediction accuracy versus total computation time for the SARCOS data. A squared exponential covariance function was used. The PIC approach produced negative predictive variances, so its NLPD could not be computed.

5.7 Discussion

Based on the numerical comparison results for the three spatial datasets and the two high dimensional datasets, we can summarize the benefits of the proposed approach as three-fold. First, it provides competitive or better mean prediction accuracy for both the spatial datasets and the high dimensional datasets. As evidenced by the MSE comparisons, for the three spatial datasets, the mean predictions of the proposed approach were either more accurate than or at least comparable to those of the state-of-the-art approaches. For the two high dimensional datasets, its mean prediction accuracy was comparable to those of the state-of-the-art approaches.

Second, as evidenced by the NLPD comparisons, the predictive distribution estimated by the proposed approach provides a better fit to test data. This implies that the proposed approach provides better predictive mean and variance estimates when taken together. We designed a simple experiment to investigate this. We used the TCO dataset to compare the degrees of fitness of the predictive distributions from our approach and the PIC, RBCM, and PGP approaches when applied to the test data. In the comparison, the tuning parameters of each compared method were chosen so that the computation time was around 300 seconds, which resulted in approximately the best MSE value for each of the methods. Using these parameters, we calculated the predictive means $\hat{\mu}(x)$ and predictive standard deviations $\hat{\sigma}(x)$ for x in 4,833 test inputs, and we counted the fractions of 4,833 test observations that fell within $\hat{\mu}(x) \pm c\hat{\sigma}(x)$ for different c values, and the fractions were compared with ground truth $P(|X| \leq c)$ where $X \sim \mathcal{N}(0,1)$. The fractional numbers closer to the ground truth are better. Table 1 shows the fractional numbers for different c values. The fractions for our method were very close to the ground truth for all choices c. The PGP method has much higher fractions than the ground truth, which implies that the PGP tends to overestimate the predictive standard deviation. Both the PIC and the RBCM methods

have much lower fractional numbers than the ground truth, which implies that these two local methods significantly underestimate the predictive standard deviation.

Last but not least, the proposed patchwork kriging advances the PGP method by broadening the applicability to high dimensional regression problems, while the PGP method is practically limited to spatial regression with only two input dimensions.

c	0.5	1.0	1.5	2.0	2.5	3.0
$P(\|X\| \le c)$, $X \sim \mathcal{N}(0,1)$	0.3829	0.6827	0.8664	0.9545	0.9876	0.9973
Our method	0.4471	0.7366	0.8814	0.9487	0.9766	0.9888
PGP	0.7118	0.8850	0.9549	0.9810	0.9905	0.9944
PIC	0.0126	0.0223	0.0362	0.0474	0.0604	0.0741
RBCM	0.1057	0.2090	0.3095	0.3987	0.4809	0.5535

Table 1: Percentages of test data ranging in between the estimated predictive mean $\pm c$ the estimated predictive standard deviation. The percentages were compared with $P(|X| \le c)$ where X is a standard normal random variable. The percentage numbers closer to $P(|X| \le c)$ are better.

6. Conclusion

We presented a new approach to efficiently solve the Gaussian process regression problem for large data. The approach first performs a spatial partitioning of a regression domain into multiple local regions and then assumes a local GP model for each local region. The local GP models are assumed *a priori* independent. However, *a posterior* dependence and related continuity constraints between the local GP models in neighboring regions are achieved by defining an auxiliary process that represents the difference between the local GP models on the boundary of the neighboring regions. By defining zero-valued pseudo observations of the auxiliary process and augmenting the actual data with the pseudo observations, we in essence force the two local GP models to have the same posterior predictive distributions at the collection of boundary points. The proposed idea of enforcing the local models to have the same boundary predictions via pseudo observations is entirely different from that of (Park and Huang, 2016), creating an entirely new framework for patching local GP models. The new approach provides significantly better prediction variance accuracy than the approach of (Park and Huang, 2016), while providing computation efficiency and mean prediction accuracy that are at least comparable and sometimes better. In addition, the spatial partitioning scheme proposed as a part of the new approach makes the new approach applicable for high dimensional regression settings, while the approach of (Park and Huang, 2016) is only applicable for one or two dimensional problems. Another advantage of the proposed approach is that its prediction accuracy does not depend strongly on the choice of tuning parameters, so one can simply fine-tune the tuning parameters to minimize the total computation time. Those advantages were numerically demonstrated with five well designed numerical experiments using five real datasets featuring different patterns and dimensions. The new approach has shown better trade-offs between total computation times and prediction accuracy than the approach of (Park and Huang, 2016) and other

local-based approaches for GP regression. We believe that the proposed patchwork kriging approach is an attractive alternative for large-scale and high dimensional GP regression problems.

Acknowledgments

The authors are thankful for generous support of this work. Chiwoo Park was supported by grants from the National Science Foundation (CMMI-1334012) and the Air Force Office of Scientific Research (FA9550-13-1-0075 and FA9550-16-1-0110).

References

Tao Chen and Jianghong Ren. Bagging for Gaussian process regression. *Neurocomputing*, 72(7):1605–1610, 2009.

Marc Peter Deisenroth and Jun Wei Ng. Distributed Gaussian processes. In *Proceedings of the 32 nd International Conference on Machine Learning*, pages 1–10, 2015.

Reinhard Furrer, Marc G Genton, and Douglas Nychka. Covariance tapering for interpolation of large spatial datasets. *Journal of Computational and Graphical Statistics*, 15(3), 2006.

Tilmann Gneiting. Compactly supported correlation functions. *Journal of Multivariate Analysis*, 83(2):493–508, 2002.

Robert B Gramacy and Daniel W Apley. Local Gaussian process approximation for large computer experiments. *Journal of Computational and Graphical Statistics*, (just-accepted):1–28, 2014.

Robert B Gramacy and Herbert KH Lee. Bayesian treed Gaussian process models with an application to computer modeling. *Journal of the American Statistical Association*, 103 (483), 2008.

Cari G Kaufman, Mark J Schervish, and Douglas W Nychka. Covariance tapering for likelihood-based estimation in large spatial data sets. *Journal of the American Statistical Association*, 103(484):1545–1555, 2008.

Finn Lindgren, Håvard Rue, and Johan Lindström. An explicit link between gaussian fields and gaussian markov random fields: the stochastic partial differential equation approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(4):423–498, 2011.

Brian McFee and Gert RG Lanckriet. Large-scale music similarity search with spatial trees. In *ISMIR*, pages 55–60, 2011.

Duy Nguyen-Tuong, Jan R Peters, and Matthias Seeger. Local Gaussian process regression for real time online model learning. In *Advances in Neural Information Processing Systems*, pages 1193–1200, 2009.

Chiwoo Park and Jianhua Z. Huang. Efficient computation of gaussian process regression for large spatial data sets by patching local gaussian processes. *Journal of Machine Learning Research*, 17(174):1–29, 2016. URL `http://jmlr.org/papers/v17/15-327.html`.

Chiwoo Park, Jianhua Z. Huang, and Yu Ding. Domain decomposition approach for fast Gaussian process regression of large spatial datasets. *Journal of Machine Learning Research*, 12:1697–1728, May 2011.

C. E. Rasmussen and Z. Ghahramani. Infinite mixtures of Gaussian process experts. In *Advances in Neural Information Processing Systems 14*, pages 881–888. MIT Press, 2002.

C.E. Rasmussen and C.K.I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 2005.

Matthias Seeger, Christopher K. I. Williams, and Neil D. Lawrence. Fast forward selection to speed up sparse Gaussian process regression. In *International Workshop on Artificial Intelligence and Statistics 9*. Society for Artificial Intelligence and Statistics, 2003.

Yirong Shen, Andrew Ng, and Matthias Seeger. Fast gaussian process regression using kd-trees. In *Proceedings of the 19th Annual Conference on Neural Information Processing Systems*, number EPFL-CONF-161316, 2006.

Edward Snelson and Zoubin Ghahramani. Sparse Gaussian processes using pseudo-inputs. In *Advances in Neural Information Processing Systems 18*, pages 1257–1264. MIT Press, 2006.

Edward Snelson and Zoubin Ghahramani. Local and global sparse Gaussian process approximations. In *International Conference on Artifical Intelligence and Statistics 11*, pages 524–531. Society for Artificial Intelligence and Statistics, 2007.

Volker Tresp. A bayesian committee machine. *Neural Computation*, 12(11):2719–2741, 2000.

Jarno Vanhatalo and Aki Vehtari. Modelling local and global phenomena with sparse Gaussian processes. In *the 24th Conference on Uncertainty in Artificial Intelligence*, page 571578. Association for Uncertainty in Artificial Intelligence, 2008.